

A SIMULATION ANALYSIS OF MULTICASTING IN DELAY TOLERANT NETWORKS

Muhammad Abdulla
Robert Simon

Department of Computer Science
George Mason University
Fairfax, VA 22030, U.S.A

ABSTRACT

Delay tolerant networks (DTNs) are a class of systems that experience frequent and long-duration partitions. As in all distributed systems, DTN multicasting is a desirable feature for applications where some form of group communication is needed. The topological impairments experienced within a DTN pose unique challenges for designing effective DTN multicasting protocols. In this paper, we examine multicasting in DTNs. Unlike earlier work we assume no knowledge of node connectivity or mobility patterns. We propose the use of both single-copy and multi-copy routing DTN routing algorithms. We also explore the use of gossiping and core nodes in DTNs to decrease the number of redundant messages while maintaining high message delivery ratios. We have performed extensive evaluations of our proposed methods. Our results show that with careful protocol parameter selection it is possible to achieve high delivery rates for various system scenarios.

1 INTRODUCTION

Delay tolerant networks (DTNs) are a class of emerging systems that experience frequent and long duration partitions (Fall 2003). In DTNs an end-to-end path between the source and the destination may only exist for brief and unpredictable periods of time. With the increased use of wireless mobile devices, many new network applications fall into this category, such as wildlife tracking, military networks, and disaster recovery and emergency response systems.

In this paper, we examine the problem of multicasting in delay tolerant networks. Multicasting is the transmission of messages or packets to a group of hosts identified by a single destination address, the group id. As in all distributed system, multicasting is desirable within DTNs for applications where close coordination or collaboration between participating members is necessary. For example, sensors deployed in a military field for intrusion detection may need

to communicate with each other for complete information regarding an intruding object. In an emergency response scenario, rescue workers want to disseminate information regarding local condition and hazard levels. While such group communication requirements can be fulfilled with separate unicast operations, path availability, power and storage restrictions and application-level delivery needs necessitate the development of efficient group communication support in DTNs.

Although multicasting, or routing in general, has been studied extensively Mobile Ad Hoc networks (MANETs), multicast routing in DTNs is a challenging problem. Since an end-to-end path may not exist, more traditional MANET proactive and reactive routing schemes fail to work. Proactive routing schemes, where nodes try to keep up to date routing information for other nodes, may fail to converge while simultaneously producing high numbers of periodic update packets. In reactive routing schemes, where routing information is obtained on demand, nodes may fail to find a path to the destination. However, this does not mean that the packets cannot be delivered to the destination. Due to node mobility, different links come up and down over time, enabling nodes to achieve *eventual delivery* through a *store-and-forward* approach, which uses buffers to hold the message until the next link comes up in the end-to-end path. A necessary condition for this approach to work is the existence of an end-to-end path between source and destination in a combined connectivity graph formed by overlapping connectivity graphs over a time interval.

In general, *single-copy* or *multi-copy* routing schemes are proposed for message delivery in DTNs. In single-copy routing schemes, only a copy of the message is transferred to achieve delivery (Spyropoulos, Psounis, and Raghavendra 2004). In multi-copy schemes, more than one copy of the message are sent (Spyropoulos, Psounis, and Raghavendra 2005). Single-copy schemes tend to be more efficient in terms of traffic overhead, but delivery ratios can be low and message transfer latencies tend to be high. Multi-copy schemes try to improve delivery ratios and latencies

by sending multiple copies of the same message along multiple paths to increase the likelihood of delivery. We will examine both approaches to multicast routing.

A number of DTN routing schemes assume prior knowledge of node mobility and connectivity, or oracles, to perform message transfers (Jain, Fall, and Patra 2004; Zhao, Ammar, and Zegura 2005). In contrast, our approach assumes no knowledge about contact histories or mobility patterns. We propose several multicast routing schemes, incorporating some previous ideas in core-based routing and DTN unicasting, as well as gossiping and epidemic approaches. We also present the results of an extensive set of ns-2 based simulation experiments that compare our various approaches.

The rest of the paper is structured as follows. Section 2 goes over related work. Section 3 describes our proposed approaches for multicasting in DTNs. Section 4 describes performance results of the proposed approaches. Finally, Section 5 concludes the paper.

2 RELATED WORK

Because of frequent network partitions many of the more traditional routing techniques for MANETS will not work properly (Burleigh and Fall 2003; Fall 2003; Jain, Fall, and Patra 2004). This fact has led to recent interest in developing new approaches for routing in a DTN environment. Here we review some of this work. Since there has been little research to date on multicasting in a DTN environment, our discussion primarily focuses on multicast DTN routing.

In the *Data Mule* approach proposed in Shah et al. (2003) a number of mobile nodes perform *random* walks to collect packets, buffer them, and deliver them to wired access points. Zhao, Ammar, and Zegura (2004) introduce a route planning strategy using message ferries that travel on a trajectory to provide communication services. Either the message ferries choose a trajectory to contact nodes, or the nodes can move near to pre-defined trajectory at a certain time to exchange packets. This type of work assumes some level of knowledge regarding node mobility and connectivity.

Jain et al formulates the DTN routing in terms of a directed multi-graph, where more than one edge may exist between a pair of nodes (Jain, Fall, and Patra 2004). Such multiple edges exist because there may be more than one distinct physical connection or different network links may only be available at different time intervals. By using different levels of information regarding connectivity and/or mobility, routing decisions can be made at individual nodes.

Other recent DTN routing approaches concentrate on trading off message complexity versus increasing the likelihood of message delivery. To limit the number of messages present *single copy* routing schemes allow only one copy of the message at a time to be present in the network. *Direct*

transmission is the simplest form of single copy routing, where each source node keeps its messages until it comes into direct contact with the respective destination nodes. Under this scheme only one message transfer is made per delivered message, incurring minimal message passing. However, in intermittently connected networks, such an approach may produce low delivery ratios and has an unbounded delivery delay (Grossglauser and Tse 2001). An improved scheme is *randomized routing*. In randomized routing, a node *A* hands over a message to another node *B* with probability $p > 0$. However, the progress of the message towards the destination can be marginal unless contact information is utilized to make routing decisions, as in *utility-based routing* or other hybrid approaches (Spyropoulos, Psounis, and Raghavendra 2004).

Generally, single copy schemes are more efficient in terms of reducing traffic overhead. However, message delivery ratios are normally lower while delivery delays are high. One way to improve delivery performance is to use multiple copies of the same message within the network. Each copy can take a different path, thereby increasing the likelihood of delivery as well as decreasing the delay. A variant of this basic approach is to allow each copy to be divided into multiple chunks using techniques such as erasure coding (Wang et al. 2005). These approaches allow multiple messages to be reconstructed at the destination.

One policy to implement a multi-copy scheme is to use simple flooding. However, due to frequent network partitions and excessive overhead a better approach is to use Epidemic Routing (Vahdat and Becker 2000). In Epidemic Routing when a pair of nodes comes into contact the nodes exchange any missing packets. Given enough storage space, Epidemic Routing can be used to reliably disseminate data across the network. By keeping a history of past encounters, nodes can reduce the overhead of Epidemic Routing (Lindgren, Doria, and Scheln 2003). However, due to its large overhead, a flooding scheme such as Epidemic Routing may not be applicable under circumstances where storage and power supplies are limited.

To address overhead problems in flooding, different forms of controlled flooding have been proposed. For instance, Spyropoulos, Psounis, and Raghavendra (2005) present Spray-and-Wait. In this method, a total of L copies of a message are initially spread to other “relay” nodes. If the destination is not found in this phase, each of the nodes carrying a copy of the message will perform direct transmission. In essence, Spray and Wait is a type of controlled flood. No mobility or connectivity information regarding the nodes in the network are assumed to be known for this scheme to work. In our work, we extend spray-and-wait to multicasting in DTNs.

Multicasting has been studied extensively in MANETS. Typical approaches for MANET multicasting include MAODV, a multicast version of the well-known AODV

MANET routing algorithm (Royer and Perkins 1999), and ODMRP, which uses a hybrid mesh and tree based approach (Lee, Su, and Gerla 2000). However, most MANET multicast routing techniques essentially assume that network partitions are relatively rare and short lived, which is not the case for DTN networks.

As mentioned above there has been relatively little work done on DTN multicasting. One approach to multicasting is simply to use multiple unicasts. This method may be difficult to achieve in practice, especially in light of frequent network partitions and uncertainty in determining the set of current group members. The problem of DTN group membership semantics is studied by Zhao, Ammar, and Zegura (2005); they also propose and evaluate several DTN multicasting policies, including a broadcast-based approach, which is basically network flooding, and *group-based* routing using the idea of multicast forwarding group (FG) node. The FG node idea is borrowed from early work in MANET Multicasting (Chiang, Gerla, and Zhang 1998).

Unlike the above work in this paper we study multicasting in DTNs when there is no available knowledge or exchanged knowledge regarding node mobility, contacts and connectivity. We use a temporal membership model for multicast delivery semantics described in Zhao, Ammar, and Zegura (2005). Specifically, we assume that each transmitted packet is stamped with a group address and a membership interval. If a node is a member of the group at any moment during the membership interval then that node is considered to be a member of the group and the message is delivered to the node.

3 MULTICASTING IN DTNS

In principle DTN multicasting can involve either a single or multiple copy approach. In this section we examine both techniques and propose a set of DTN routing policies, including multicast direct transmission, a multicast variant of Spray and Wait, and multicast versions of epidemic routing and gossiping. We also propose the use of special *core based* nodes as a method to control buffer space usage. The remainder of this section describes these policies, and the following section shows the results of our performance analysis.

3.1 Direct Transmission

Direct transmission can have different forms for multicasting. The simplest approach is that the sender will try to deliver its packet to every group member directly. This kind of source based delivery (SBD) is similar to using multiple unicasts to achieve multicasting, and is not really considered a form of multicast.

Instead, we implement direct transmission within a group. In this scheme, group members try to deliver packets

for that group among each other. Members of a group do not give packets to nodes that do not belong to the group, nor do they take packets from them. Members of the same group exchange packets when they come into direct contact with other. This scheme is called *group-based direct delivery*.

Another alternative is to allow different group members to exchange packets. In this scheme, members of a group transfer packets with nodes that are members of other groups. This, however, normally requires multiple copies of a packet to be sent out.

3.2 Epidemic Routing

In terms of message overhead the opposite approach from direct transmission is flooding. In a flooding approach every node that receives a packet broadcasts it to all of its neighbors. However, in intermittently connected networks simply broadcasting packets may not achieve the goal of reaching as many nodes as possible due to network partitions. In such a context, an *Epidemic Routing* approach outperforms flooding (Vahdat and Becker 2000).

Figure 1 shows how two nodes can exchange messages in Epidemic Routing. When two nodes come into contact, each node will exchange message information to see if there are any messages that the other node has that it has not received. Message indexes are sent as a summary vector. After such pair-wise exchange of messages, each node will get all the messages carried by the other node that it has not received by far. This means that as long as buffer space is available, messages spread like an epidemic disease among nodes through “infection”.

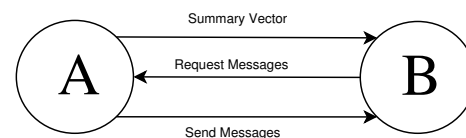


Figure 1: Message Exchange in Epidemic Routing

To determine which packets have been previously seen, there must be a globally unique message ID. We propose the use of a tuple $(source_id, sequence_number)$ where $source_id$ is the id of the sending node and $sequence_number$ is a unique sequence number for each message sent by the node. A time limit is used as the minimum time span between two exchanges for any given pair of nodes to reduce the number of vector exchanges. For a large number of messages, the summary vector can become quite large. To reduce the size of the summary vector, we are investigating the use of a *Bloom Filter* technique.

3.3 Spray and Wait

Although epidemic routing can achieve high delivery ratios and low delays, it requires that nodes to have sufficiently large buffer spaces and can incur large message transfer overhead. An alternative approach is to control the level of flooding using techniques such as *Spray and Wait* (Spyropoulos, Psounis, and Raghavendra 2005).

In the context of unicast, Spray and Wait works as follows. L number of copies are initially spread over the network by the source or other nodes to L distinct relays (spray phase). If the destination is not found during this phase, each node that is carrying a copy of the message performs a direct transmission.

In group-based message communications, we modify Spray and Wait as follows. The source node generating a message gives out copies of the message to other nodes that came into contact, similar to Spray and Wait in unicast routing. However, the destination is a group of nodes, as opposed to a single node in unicast. As a result, after delivering a message to destination group, the sending node may have to keep the copy of the message as there are other group members to deliver the message. This is different from Spray and Wait case in unicast, where the sending node immediately deletes the message copy once it delivers it to the destination. At the same time, sending out copies of a message to group members by the source or the relays till the message expires may generate considerable amount of unnecessary traffic where multiple copies of the same message are delivered to group members by different intermediate carriers.

In this approach, members of a group exchange packets within their group, as well as members of other groups. A number of issues arise in this approach, including controlling number of copies of packets, how packets are exchanged within and among groups, etc.

3.4 Gossiping

Because of problems in accurately determining group membership, nodes will have to find a way to decide when to stop sending a message to destination group members even before the message expires. To achieve this goal, in addition to a TTL a gossip based approach can be used. In this approach, nodes send packets with a probability $p < 1$, rather than always sending packets as specified by original protocol.

3.5 Core-Aided Routing

Some DTN systems, such as sensor networks, may have resource constrained nodes. One approach to addressing this problem is to assume that some nodes are not resource constrained, and to designate these nodes as *core nodes*.

When regular nodes come into contact with core nodes, they can delegate all or some of their packets to the core node. When nodes meet these core nodes, they will dump their messages to the cores, and delete their own local copies. This can be helpful in environments where buffer storage for most nodes is very stringent. One potential drawback to this approach is a possible increase in message delivery delay. This is because fewer nodes carry message copies, thereby decreasing the chance of contacting destination nodes in a short time.

4 EXPERIMENTAL RESULTS

In this section, we present a performance study of several classes of DTN multicast algorithms. Our results were obtained from simulation experiments using $ns - 2$ and our own code. Specifically we evaluate direct transmission, direct transmission with group based Epidemic Routing (DTGER), multicast Spray and Wait (M-SW), and Multicast Spray and Wait with group based Epidemic Routing (M-SWER). We will see that adding Epidemic Routing approach to both direct transmission and multicast Spray and Wait improves performance in terms of delivery ratios and message delivery delay. We also evaluate the effectiveness of a protocol called M-SWG, which is gossiping added to M-SW. Finally, we examine how the introduction of DTN core nodes effects performance.

4.1 Metrics and Methodology

All of our experiments used one of the standard $ns - 2$ mobile wireless models, including the default transmission model which has a 250m radio range. We collected statistics for average message delivery rate, full delivery ratio, and the average message delivery delay are used as the performance evaluation metrics.

The *average message delivery ratio* is the ratio of delivered messages to the number of messages that should have been delivered to destination group members according to the multicast delivery semantics being used. Because each message to a group should be delivered to all group members, we use the total number of messages to be delivered rather than total number of messages generated to normalize the ratio. The average message delivery ratio reflects the overall efficiency of the method in delivering messages.

In many types of multicast applications, however, it is important to deliver the message to all the group members. For this, we use the term *full delivery ratio* or *total delivery ratio*. The full delivery ratio is the ratio of the total number of messages delivered to all of the members of the target group to the total number of generated messages. The full delivery ratio reflects the how often the method successfully delivers each message to all group members. Therefore, we record failure if at least one group member does not

receive a message. The *average delivery delay* is the average delay of all the messages delivered to the destination. The delay of a delivered message is calculated by subtracting the delivery time by the message generation time. Finally, to evaluate M-SWG we define the *goodput* to be ratio of total number of messages deliveries to the total number of messages transferred.

The default settings in our simulations are as follows. Each simulation run has 40 nodes in a $6000m \times 6000m$ area. Each node sends a message to its group at every 100 seconds on average (the time between two message generations is chosen randomly from 0 to 200 seconds). The *HELLO* messages, which act as *heart beat* or *beacon* messages, are generated every 3 seconds. Each node has a storage space that can hold 500 messages. Messages expire after 5000 seconds and will not be delivered if expired.

The simulations use the random-waypoint (RWP) mobility model. In RWP nodes randomly choose a point in the area and moves towards that destination with an average speed uniformly distributed between $speed_{min}$ and $speed_{max}$. In our simulations, $speed_{min}$ is 3, $speed_{max}$ is 10. Pause time after reaching the destination point is 3 seconds.

We run each experiment with a random seed for at least 20 times. The simulation time is 40,000 seconds, after which most results are found to be stabilized.


4.2 General Scalability Results

Figure 2 shows the results of a scalability test for M-SW. In these experiments, the size of the simulation area is changed from $3000m \times 3000m$ to $7000m \times 7000m$, increasing the square shaped area by $1000m$ each time. The number of nodes is changed from 20 to 70, increasing by 10.


Subfigures (a) and (c) show that average delivery ratios tend to be higher and average delay of delivered message tend to be lower in denser networks. Subfigure (b) demonstrates a similar trend for the full delivery ratio. However, in settings where large numbers of nodes are distributed in larger areas full delivery ratios are lower. In such cases, even though increasing the number of nodes increase network density, the average number of group members also increases, which decreases the likelihood of a message being delivered to all of the members of the group.

4.3 Performance Comparisons


We first do a performance comparison of direct transmission, DTGER, M-SW and M-SWER, shown in Figure 3. As can be seen in figure, the M-SW approach has a higher delivery ratios. Adding inter-group Epidemic Routing helps direct routing more than it helps Spray and Wait. This can be explained by the fact that the basic M-SW approach already has the mechanism to exchange packets among a node and its neighbors.



(a) Delivery Ratio



(b) Full Delivery Ratio



(c) Delivery Delay

Figure 2: Performance Characteristics of Multicast Spray and Wait

As shown in Figure 4, the full delivery ratio of direct transmission is very low, nearing zero. Compared to average delivery ratio differences shown in Figure 3, M-SW shows a substantial improvement in the full delivery ratio when inter-group Epidemic Routing is added. In short, for either direct and M-SW, adding an Epidemic Routing approach improves full delivery delivery ratio more than the average delivery ratio.

Figure 5 shows delivery delays. M-SW has lower delay than direct transmission. Adding inter-group Epidemic

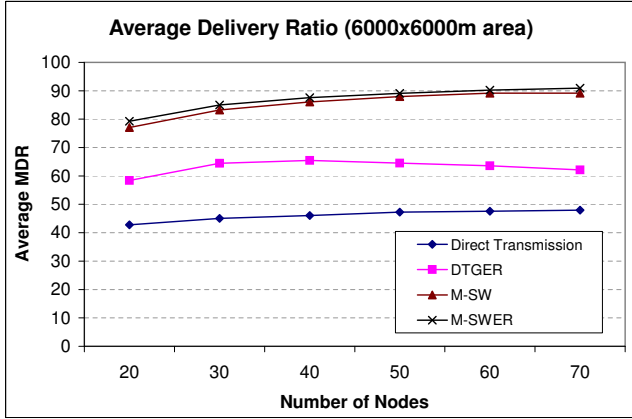


Figure 3: Comparison of Average Delivery Ratios




Figure 5: Comparison of Delivery Delays

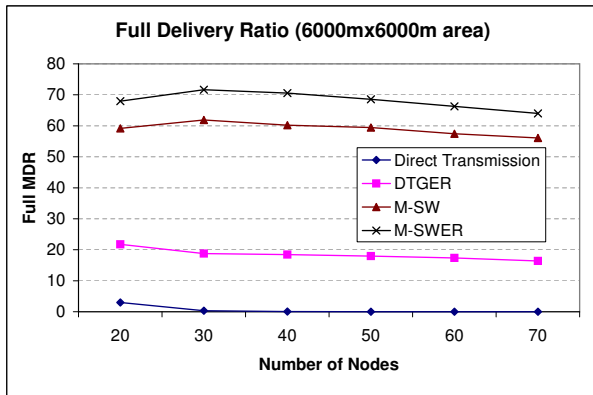


Figure 4: Comparison of Full Delivery Ratios




Figure 6: Effects of Gossip Probability

Routing lowers the delivery delay for both M-SW and direct transmission.

4.4 Effects of Gossiping on Performance

We next show the results for M-SWG. Under the M-SWG policy each node sends a message to a contact with a probability p , where $0 < p < 1$. If the probability is 1, then it becomes M-SW without gossip. Figure 6 shows the changes in average message delivery ratio, full delivery ratio, delay, goodput, and total number of packets sent when p is changed from 0.6 to 1. For comparison, each series of values are normalized with the value when $p = 1$.

As depicted in Figure 6, the total number of messages decreases as gossip probability p decreases. The average message delivery ratio decreases more slowly than total number of packet sent. This means the goodput increases as we decrease p , which is also shown in the graph. The

delivery delay also shows a slow change relative to the total number of packets. The full delivery ratio, however, decreases at a faster rate than the number of packets. These results show the tradeoffs of using a gossip approach in different network scenarios. When network resources such as power and buffer space are scarce and the full delivery ratios are not very important in the application, one can sacrifice slight levels of average delivery ratio to achieve a noticeable increase in goodput.

4.5 Core Nodes

In some application scenarios where nodes have small buffers compared to application requirements, it may be helpful to introduce some nodes with large buffer sizes to act as core nodes or super nodes. In our approach, on contacting a core node a regular node dumps all of its packets to the core if its buffer space is above some threshold. This can help alleviate buffer constraints on regular nodes and decrease the number of messages dropped due to buffer shortage. Figure 7 shows the effects of introducing core nodes in a system




Figure 7: Effects of Core Nodes on Delivery Ratios

where regular nodes have very small buffer spaces. In this particular configuration, 40 nodes are simulated moving in a $6000m \times 6000m$ area. Each regular node has a buffer size of 50 packets. Core nodes with a buffer size of 2000 packets each are introduced in different ratios and their effects to delivery ratios are studied.

We see that in many scenarios the introduction of core nodes improves overall system performance. In fact, performance is improved even if we do not use the regular-to-core transfer policy but only add core nodes with larger buffer sizes. We show this case separately (denoted as “no ast.” in the graph) to distinguish the effects.

Figure 7 describes the changes in the message delivery ratio and the full delivery ratio when the ratio of core nodes is changed from 0 (no core node) to 1/3. We observe that the delivery performance improves as more core nodes are introduced. Using the core assisted approach produces higher delivery ratios, especially full delivery ratios. However, the average delivery delay also increases with the introduction of core nodes. This is because even though more messages get delivered to the destination those messages delivered by the core nodes those same messages reside for longer periods of time in the cores.

One important observation is that the full delivery ratio is very sensitive to parameter changes. A small increase in average delivery ratio can result in a considerable improvement in the full delivery ratio. At the same time, an effort to reduce redundant message transfers by trading off small decreases in average delivery ratio can cause a degradation in the full delivery performance. When the full delivery ratio is an important factor in an application, these performance indicators must be taken into account.

5 CONCLUSIONS

This paper studied the problem of multicasting in delay tolerant networks (DTNs). We gave an overview of related problems and the current research in this category of networks. Under circumstances where node mobility and contact information are not available, DTNs presents unique challenges for routing approaches because of long delays and frequent partitions. We proposed and evaluated several routing algorithms for DTN multicasting.

Based on our simulations, we obtained the following results. Generally, multi-copy routing schemes performs better in terms of delivery ratios and delivery delays. Intra-group multi-copy schemes can be combined with inter-group epidemic schemes to obtain better performance, especially for full delivery ratios. Gossiping provides a good mechanism mechanism for trading off goodput and delivery ratios. Utilizing gossip can lead to higher goodput ratios with a slight decrease in average delivery ratio, but also heavily decreases the full delivery ratio. In general, it is necessary to have a considerable amount of redundancy to achieve modest levels of full delivery ratio. Finally, for the case where many nodes have small buffer spaces, we found that introducing high-buffer capacity core nodes increases delivery ratios.

REFERENCES

- Burleigh, S., and K. Fall. 2003, June. Delay tolerant networking: An approach for interplanetary internet. *IEEE Communications Magazine*.
- Chiang, C.-C., M. Gerla, and L. Zhang. 1998. Forwarding group multicast protocol (FGMP) for multihop, mobile wireless networks. *Cluster Computing* 1 (2): 187–196.
- Fall, K. 2003, August. A delay-tolerant network architecture for challenged internets. 27–34. In *Proceedings of ACM SIGCOMM'03*. Karlsruhe, Germany.
- Grossglauser, M., and D. N. C. Tse. 2001. Mobility increases the capacity of ad-hoc wireless networks. In *INFOCOM*, 1360–1369.
- Jain, S., K. Fall, and R. Patra. 2004, August. Routing in a delay tolerant network. In *Proceedings of ACM SIGCOMM'04*. Portland, Oregon.
- Lee, S., W. Su, and M. Gerla. 2000. On-demand multicast routing protocol in multihop wireless mobile networks. *ACM/Baltzer Mobile Networks and Applications, Special Issue on Multipoint Communication in Wireless Mobile Networks*.
- Lindgren, A., A. Doria, and O. Scheln. 2003. Probabilistic routing in intermittently connected networks. In *Proceedings of the Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2003)*.

- Royer, E., and C. Perkins. 1999, August. Multicast operation of the ad-hoc on-demand distance vector routing protocol. In *Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom)*, 207–218. Seattle, WA.
- Shah, R., S. Roy, S. Jain, and W. Brunette. 2003, May. Data mules: Modeling a three-tier architecture for sparse sensor networks. In *IEEE Workshop on Sensor Network Protocols and Applications (SNPA)*.
- Spyropoulos, T., K. Psounis, and C. Raghavendra. 2005, August. Spray and wait: An efficient routing scheme for intermittently connected mobile networks. In *Proceedings of the ACM SIGCOMM'05 Workshop on Delay-Tolerant Networking*. Philadelphia, PA.
- Spyropoulos, T., K. Psounis, and C. S. Raghavendra. 2004, October. Single-copy routing in intermittently connected mobile networks. In *Sensor and Ad Hoc Communications and Networks*, 235–244. IEEE.
- Vahdat, A., and D. Becker. 2000, April. Epidemic routing for partially connected ad hoc networks. Technical Report CS-200006, Duke University.
- Wang, Y., S. Jain, M. Martonosi, and K. Fall. 2005, August. Erasure-coding based routing for opportunistic networks. In *Proceedings of the ACM SIGCOMM'05 Workshop on Delay-Tolerant Networking*. Philadelphia, PA.
- Zhao, W., M. Ammar, and E. Zegura. 2004, May. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proceedings of the Fifth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2004)*.
- Zhao, W., M. Ammar, and E. Zegura. 2005, August. Multicasting in delay tolerant networks: Semantic models and routing algorithms. In *Proceedings of the ACM SIGCOMM'05 Workshop on Delay-Tolerant Networking*, 268–275. Philadelphia, PA.

AUTHOR BIOGRAPHIES

MUHAMMAD ABDULLA is currently a Ph.D. student in the Computer Science Department of George Mason University. His research interests are in the field of sensor networks and delay tolerant networks. His email address is <mabdulla@cs.gmu.edu>.

ROBERT SIMON is an associate professor in the Computer Science Department of George Mason University. He received his Ph.D. in Computer Science from the University of Pittsburgh. His research interests are in networking, distributed systems and simulation. He has published more than 70 papers on these topics. He serves on numerous program committees, and was the Program Co-Chair for IEEE MASCOTS 2006. His email address is <simon@cs.gmu.edu>.