Delay-Differentiated Gossiping in Delay Tolerant Networks

Parameswaran Ramanathan and Aarti Singh
Department of Electrical and Computer Engineering
University of Wisconsin-Madison
parmesh@ece.wisc.edu, singh@cae.wisc.edu

***Abstract*— Delay Tolerant Networks are increasingly being envisioned for a wide range of applications. Many of these applications need support for quality of service (QoS) differentiation from the network. This paper proposes a method for providing probabilistic delay assurances in DTNs.**

In particular, the paper presents a method called Delay-Differentiated Gossiping to assure a certain probability of meeting the packets' delay requirements while using as little network resources as possible. The idea is to adapt a set of forwarding probabilities and time-to-live parameters to control the usage of network resources based on how the delay requirements are being met. Empirical results evaluating the effectiveness of the proposed method are also included. The results show that there are simple ways of assuring the delay requirements while making effective use of the network resources.

I. Introduction

Delay Tolerant Networks (DTNs) are based on the premise that the network is often highly partitioned and there is no connected end-to-end path between most pairs of nodes at any given time instant. Such networks are increasingly being envisioned for many applications.

For example, consider a network of sensing devices carried by troops or vehicles in a battlefield. As the troops and vehicles move, the sensing devices can sample the environment at different locations, exchange their samples with other devices, and collaboratively make decisions about the environment. For instance, whether an unauthorized target is present in a given region, what is the location of a known target, and what is the current boundary of an environmental contamination [15]. The underlying network over which this collaboration occurs is a DTN if the communication range of the sensing devices is small relative to the average distances between individual troops and vehicles. In this case, the needed information exchange must occur during the occasional contacts between the devices.

Other applications envisioned for DTNs include disseminating weather and traveler information in a national park through devices carried by hikers [1], providing infrastructure access to isolated village through communication equipped rural buses [12], and wildlife monitoring through devices attached to animals under study [7]. Several other networks such as ad hoc network of satellites [13] and Inter-planetary networks [2] can also be viewed as a DTN.

Recent research has shown that conventional network protocols do not work well in DTNs [3], [4]. For instance, Transmission Control Protocol (TCP) and Internet Protocol (IP) based applications implicitly assume that there exists an end-to-end path between the two communicating nodes and the round-trip communication delay between two nodes is not excessive. Since these two assumptions do not usually hold in highly partitioned networks, the TCP/IP based protocols are not well suited for DTNs. Consequently, there is growing interest in developing new protocols tailored for DTNs [3], [4], [6], [9], [10], [16].

For example, an approach called message ferrying is proposed in [16], [17], where special mobile nodes called message ferries help in carrying information from one stationary node to another. The movement of the ferries are planned to improve performance metrics such as fraction of successfully delivered packets and average energy consumption. In [6], [10], space-time routing algorithms are proposed for minimizing the delay of message delivery in networks with predictable mobility. These algorithms exploit knowledge of node locations in near future to plan a carry and relay strategy to deliver messages from one node to the other. Reliable information delivery through a mechanism called custody transfer is considered in [4]. The basic idea is to transfer the responsibility for reliable delivery from one node to the other until the information is successfully delivered to the destination. A probabilistic routing scheme for DTNs is proposed in [9] where nodes use the history of meetings between them to estimate the probability of successful deliveries to each destination through other nodes and thereby identify the best next hop for each possible destination.

Note that, none of the existing work on DTNs consider the problem of providing quality of service (QoS) assurances to ongoing communications. However, such assurances are important for many DTN applications. Packets in the network may differ in their delay requirements, e.g., event alerts may have shorter delay requirements than data collection messages in sensing networks. Prioritization of packets is important in DTNs since each node is responsible for multiple data gathering tasks and the contacts between nodes are very infrequent. This paper describes an approach for providing delay assurances in DTNs.

In this paper, the network traffic is grouped into one of M service classes based on the delay requirements, i.e., packets in a service class are assumed to have the same delay requirements. The proposed approach strives to meet the delay requirements of each packet while making efficient use of the network resources. This is done as follows. The network maintains a set of parameters denoted by p_1, p_2, ..., p_M. A packet belonging to service class i is forwarded using a probabilistic variant of the basic gossiping scheme in which each node forwards the packet with probability p_i. The smaller the probability the larger the expected delay but smaller the amount of the network resources consumed in delivering the packet. The network adapts the probabilities based on the observed delays so as to assure a certain desired probability of meeting the packet's delay requirement.

Note that this approach is considerably different from the conventional approaches for providing delay assurances. This is because the conventional approaches are designed for non-DTN networks in which queuing delay at each node is the primary component of end-to-end delivery delays. In contrast, in DTNs, the primary component of end-to-end delivery delay is the delay in establishing a communication link between two nodes. Therefore, queuing disciplines that re-order message transmissions are not effective in supporting delay assurance. Gossiping based approaches are better suited for DTNs.

The rest of this paper is organized as follows. The system model and problem formulation are described in Section II. The proposed schemes for delay assurances in the network is described in Section III. Results of an empirical evaluation of the proposed scheme are presented in Section IV. The paper concludes with Section V.

II. System Model and Problem Formulation

The network is comprised of N nodes labeled 1, 2, ... , N. Each node has a wireless communication capability, but with a limited range R. Since nodes are mobile, the set of other nodes that are in its communication range varies with time. However, unlike conventional mobile wireless ad hoc networks, N, R, and the geographic area in which the nodes move are such that the network is highly partitioned most of the time.

One or more of the nodes are base nodes towards which all traffic in the network is directed. For instance, a base node may be a fusion center in a sensor network that aggregates the data from all the other nodes or the base nodes have access to external network infrastructure. The base nodes may be stationary or mobile nodes on vehicles with satellite access.

A node s may have packets for one of the base nodes d. Depending on the regions node s and possibly node d move, they may never come within R of each other. In that case, node s must rely on other nodes to forward packets on its behalf. The forwarding of packets in the network is based on carry and relay. That is, node s may carry the packet intended for node d until it comes in range of another node say i. Node s then relays the packet to node i, which in turn carries the packet until it meets another node say j and so on until the packet reaches node d.

Each packet has a deadline by which it must be delivered to the destination. The deadline is assumed to be one of M possible values D_1, D_2, ..., D_M. Without loss of generality assume that $D_1 > D_2 > \cdots > D_M$. For simplicity of presentation, a packet with deadline D_i is said to be in service class i. The deadlines of the packets are soft in the sense that a few deadline misses are not catastrophic to the application. The network strives to assure that the probability of a packet being delivered before its deadline exceeds a pre-specified threshold Δ.

To achieve this delay assurance, the network expends more resources to deliver a packet of shorter deadline, i.e., the network expends more resources to deliver a packet of class i than that for a packet of class j, $j < i$. The challenge is to expend as little as resources as possible while still meeting the above probabilistic deadline assurance goal.

III. Probabilistic Delay Assurances

A primitive gossiping scheme is equivalent to flooding [14] and works as follows. If a node has a packet destined for a node d, it locally broadcasts the packet to all its neighbors. The neighbors, in turn, broadcast the packet to all their neighbors, and so on. The packet eventually reaches node d. Since other nodes do not know whether or not the packet has already reached node d, they continue to forward until all nodes have received the packet. However, a node does not forward a particular packet more than once.

The key advantage of primitive gossiping is that each packet is delivered in the shortest possible time. However, there is considerable wastage of network resources. Multiple copies of each packet are delivered to the destination. Furthermore, packets continued to be forwarded in the network even after it has been delivered to the destination. Certain simple solutions can be used to reduce the resource wastage at the possible expense of increased delay in delivering the packet [8], [5]. For example, to reduce the wastage due to unnecessary transmissions after a packet has been delivered to the destination, one can include a *Time-To-Live* field in the packet. If a node receives a packet that has been in the network for a time greater than the value in the Time-To-Live field, then the packet is not forwarded. If the value for the Time-To-Live is carefully chosen based on the expected delay in delivering a packet to the destination, then this approach can reduce some of the resource wastage. Another simple modification is to randomize the scheme [5] so that all nodes, except the source node, forward a received packet with certain probability $p < 1$. Smaller p results in less wastage of network resources.

Based on these ideas, this paper proposes the following modification of primitive gossiping to support delay assurances in DTN called *Delay Differentiated Gossiping* (DDG). The solution can be described in three parts: (i) what a node does when a new packet is generated, (ii) what a node does when it receives a packet from another node, (iii) what the destination does when it receives a packet.

(i)*What a node does when a new packet is generated*: The network maintains and updates M probabilities $p_1, p_2, \ldots, p_M$ and M time-to-live parameters T_1, T_2, ..., T_M, as described below. When a node s generates a packet of class i, it forwards the packet to all the nodes it meets until a Time-To-Live duration T_i. Before forwarding the packet, it also includes the current value of p_i and T_i in the message header.

(ii)*What a node does when it receives a packet from another node*: When a node receives a packet of class i, it decides to forward the packet with probability p_i and to not forward the packet with probability $(1 - p_i)$. If the node decides to forward, it relays the packet to all the nodes it meets until the Time-To-Live field in the packet (T_i) has expired. Otherwise, it discards the packet and all of its copies.

(iii)*What the destination does when it receives a packet*: The destination base node maintains $p_1, \ldots, p_M$. It receives most of the packets; only the packets whose Time-To-Live duration expire before delivery are never received. If a packet of class i is delivered before its deadline, the destination reduces parameter p_i by a multiplicative factor β_p. Likewise, if a packet of class i is delivered after its deadline, the destination increases parameter p_i by a multiplicative factor γ_p. The parameters $p_1, \ldots, p_M$ are periodically relayed to all the nodes. Packets also have sequence numbers. From the sequence numbers, the destination can deduce that certain packets never got delivered (after some delay). When it discovers that a certain class i packet was never delivered, parameter p_i is increased by a multiplicative factor γ_p.

A similar approach is used to maintain $T_1, \ldots, T_M$. When the destination discovers that a certain class i packet was never delivered, it increases T_i by a multiplicative factor γ_T. Each time a packet is successfully delivered, it decreases T_i by a multiplicative factor β_T. The parameters $T_1, \ldots T_M$ are periodically relayed to all nodes.

The selection of parameters β_p and γ_p is based on the following observation. Given a probability of forwarding p, there is a certain probability of not assuring the packet's deadline. The probability of not assuring a packet's deadline tends to decrease monotonically with p. Let $m(p)$ denote the probability of not assuring a packet's deadline for a given probability of forwarding p. In steady state, when the probability of forwarding converges, the following relationship should be satisfied.

$$p = \gamma_p \times p \times m(p) + \beta_p \times p \times (1 - m(p)).$$

After simplification, the above relationship can be rewritten as

$$m(p) = \frac{1 - \beta_p}{\gamma_p - \beta_p}.$$

If the desired assurance threshold is Δ, then the desired value for $m(p) = 1 - \Delta$. Therefore, β_p and γ_p must be selected such that

$$1 - \Delta = \frac{1 - \beta_p}{\gamma_p - \beta_p}.$$

There are many values of β_p and γ_p that satisfy the above relationship. The larger the difference between γ_p and β_p, the larger the variability in the probability of forwarding but faster the convergence.

A similar relationship can be derived for selecting β_T and γ_T. If η denotes the maximum tolerable fraction of packets not delivered to the destination before Time-To-Live duration expires, then

$$1 - \eta = \frac{1 - \beta_T}{\gamma_T - \beta_T}.$$

Here again, there are many values of β_T and γ_T that satisfy the above relationship. The larger the difference between γ_T and β_T, the larger the variability in the Time-To-Live duration but faster the convergence of the Time-to-Live duration to a steady value.

IV. Empirical Evaluation of Delay Assurances

We consider a network of nodes moving in a specified geographic area based on the following homing mobility model.

Homing Mobility model: The location of the base node is considered the home location for all nodes. A node starts at home and moves as follows. It selects, from a uniform distribution, a random location from the entire geographic area and moves to the chosen location at a constant speed. It then pauses for a randomly selected duration and then either chooses to return home or moves to another randomly selected location. The process repeats throughout the simulation.

Traffic model: Recall that, we assume that queuing is not the primary contributor to end-to-end delay. Instead, the end-to-end delay is primarily due to lack of connectivity between the nodes. In particular, when two nodes meet, we assume that the duration of their meeting is sufficiently long to transmit all the packets that need to be exchanged between them.

When a packet is generated, its source is selected uniformly at random. The class of the packet is also randomly selected with equal probability for all classes.

Performance measures: For each class, the simulator maintains the fraction of packets that were delivered to the destination after the deadline and the fraction of packets that were never delivered to the destination before its Time-To-Live duration expired. The number of packets simulated is long enough to ensure that these measures for each class are known to sufficient confidence. The simulator also collects the average number of times a packet is transmitted in the network, henceforth referred to as *average packet spread*. Packet spread is directly proportional to the network resources (bandwidth and energy) consumed in delivering a packet.

Simulation parameters: For the results presented in this section, the various simulation parameters are as follows.

Number of nodes	50
Number of base nodes	1
Geographic area	100×100 units
Radio range of a node	10 units
Probability of returning to home	0.1
Number of service classes	4
Initial Probability of gossiping	$p_1 = 0.25$, $p_2 = 0.50$, $p_3 = 0.75$, $p_4 = 1.00$

At each time instant, the simulator computes the number of nodes in the largest connected component. Figure 1 shows the

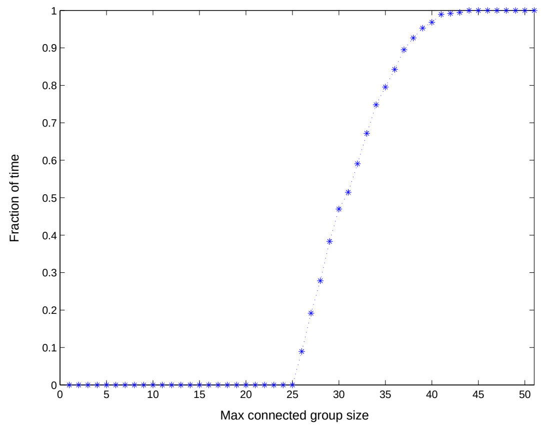

Fig. 1. Distribution function characterizing the number of nodes in the largest connected component of the simulated network.

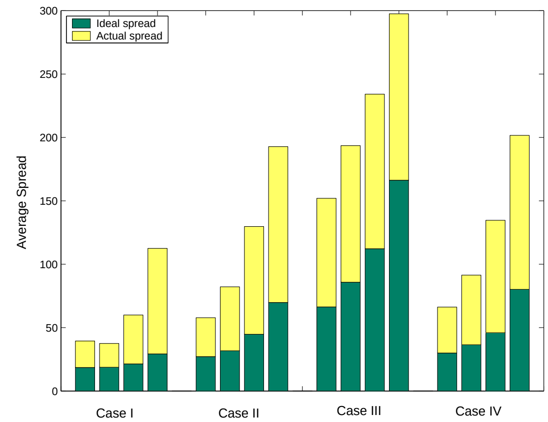

Fig. 3. Average spread for each of the four classes in a DTN using DDG scheme for four cases corresponding to different base node locations.

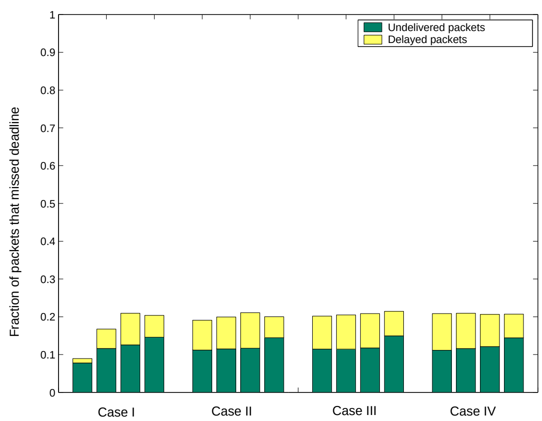

Fig. 2. Fraction of packets for each of the four classes that missed the deadline in a DTN using DDG scheme for four cases corresponding to different base node locations. The desired assurance threshold was set at 0.8.

fraction of time instances the size of the largest connected component is smaller than different values on the x-axis. Observe that a majority of the time instances, the size of the largest connected is no more than 60% of the total number of nodes. This shows that the simulation parameters selected above result in a network that is highly partitioned most of the time, i.e., the simulated network is a DTN.

Performance of DDG: A packet's deadline is considered to be assured only if it is delivered to the destination before its deadline. Packet that are either delivered late or those that are never delivered are considered to have missed their deadline. Figure 2 shows the fraction of packets in each class that missed their deadlines, and the relative split between the late and not delivered packets, for four cases corresponding to different base node locations. In Case I, the base node is located in the center of the geographic region. In Case II, the base node is located at location (25, 25). In Case III, the base node is located at a corner of the geographic region. In Case IV, the base node is also mobile and the other nodes do not return to home, they just follow the random waypoint mobility model. The desired assurance threshold is set at 0.8 for all the four cases. Observe that the fraction of missed deadlines is less than 0.2 (i.e., fraction of packets whose deadlines were assured is greater than 0.8) in all the cases. In Case I, the fraction of missed deadlines is smaller than 0.2 for some classes. This is because in a random waypoint mobility model the probability distribution of mobile node locations tends to be higher near the center of the geographic region [11], where the base node is located.

Figure 3 shows the average spread for each class in the above mentioned cases. As discussed earlier, one of the reasons for a large spread is due to the difficulty in determining when a packet has been delivered to the destination. To assess the impact of this difficulty on the average packet spread, Figure 3 highlights the average spread in an idealized case where all nodes in the network instantly become aware of the fact that the packet has been delivered to the destination and they immediately stop forwarding the packet. Observe that in each of the cases, the average spread for class 4 (rightmost bar) is substantially higher than that for class 3 which in turn is much higher than that for class 2 and class 1 (leftmost bar), respectively. This shows that the network utilizes more resources to deliver class 4 packets that have higher priority (smaller deadlines) than class 1 packets. Further observe that the average spread for a class increases from Case I to Case II to Case III. This is because the base node is centrally located in Case I, moderately centrally located in Case II and located out of the way in Case III. The ordering in the average spreads holds because the mobile nodes are more likely to be located in the center of the region than in the periphery (see discussion in the previous paragraph).

Figure 4 shows the variations in the parameters p_1, p_2, p_3, and p_4 over the course of a single simulation run. In this run, base node was stationary at location (25, 25). Observe that the probabilities tend to converge towards a steady value. Furthermore, $p_4 > p_3 > p_2 > p_1$ after convergence. This is because the deadline requirements of packets are the most stringent for class 4, followed by class 3, class 2, and class 1 respectively. The periods where the probabilities do not change

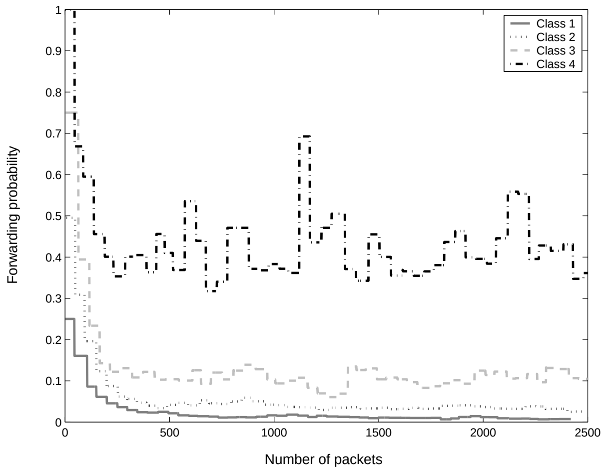

Fig. 4. Variation in p_1, p_2, p_3 and p_4 over the course of a simulation run.

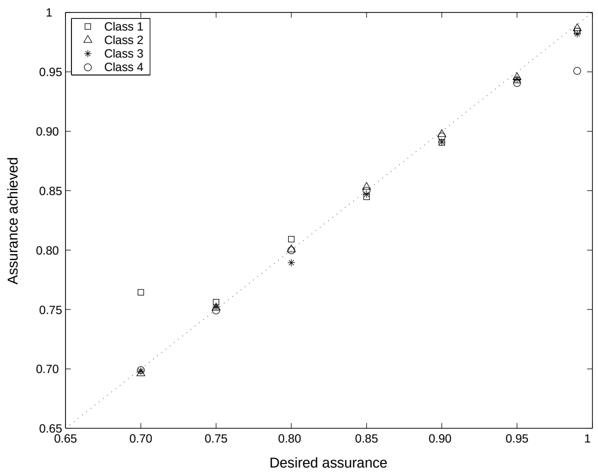

Fig. 5. Fraction of missed deadlines for different assurance thresholds.

correspond to the periodic delay in relaying the updated gossip probabilities and Time-To-Live parameters from the base node to other nodes.

Figure 5 shows the the variation in the fraction of packets whose deadlines were assured (referred to as achieved assurance) for several different desired assurance thresholds. The scheme is working well if the achieved assurance equals the desired assurance for all the classes. If the point on the graph is above (below) the line of unit slope, then the achieved assurance exceeds (lags) the desired assurance. Observe that, the achieved assurance is close to the unit slope line in almost all cases showing that the adaptation of the parameters works well in the proposed scheme. The only case where the point on the graph is significantly below the unit slope line corresponds to class 4 packets when the desired assurance is 0.99. In this case, the achieved assurance is less than the desired assurance because the network is such that even with gossiping probability of one the desired delay assurance cannot be provided.

V. Summary

Many of the applications envisioned for DTNs will need some support for quality of service differentiation from the network. In this paper, we discussed an approach for providing probabilistic delay assurances in DTNs. In this approach, the network adapts a set of design parameters to reduce the bandwidth consumed in delivering a packet before its deadline requirement. The goal is to ensure that the probability of timely delivery of packets exceeds a certain specified threshold. The simulation results in the paper show that the proposed scheme is effective in adjusting the parameters to meet the desired probability of assurance.

References

[1] A. Beaufour, M. Leopold, and P. Bonnet, “Smart-tag based data dissemination,” In *First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA02)*, June 2002.

[2] S. Burleigh, V. Cerf, R. Durst, K. Fall, A. Hooke, K. Scott, and H. Weiss, “The interplanetary internet: a communications infrastructure for mars exploration,” In *53rd International Astronautical Congress, The World Space Congress*, October 2002.

[3] K. Fall, “A delay-tolerant network architecture for challenged internets,” In *Proceedings of SIGCOMM*, pp. 27–34, August 2003.

[4] K. Fall, W. Hong, and S. Madden, “Custody transfer for reliable delivery in delay tolerant networks,” Technical Report IRB-TR-03-030, July 2003.

[5] Z. Haas, J. Halpern, and L. Li, “Gossip-based ad hoc routing,” In *INFOCOM 2002*, June 2002.

[6] S. Jain, K. Fall, and R. Patra, “Routing in delay tolerant networking,” In *Proceedigns of SIGCOMM*, August 2004.

[7] P. Juang, et al, “Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with Zebranet,” In *Proceedings of ASPLOS*, 2002.

[8] Q. Li and D. Rus, “Communication in disconnected ad-hoc networks using message relay,” *Journal of Parallel and Distributed Computing*, vol. 63, no. 1, pp. 75–86, January 2003.

[9] A. Lindgren, A. Doria, and O. Schelén, “Probabilistic routing in intermittently connected networks,” In *Proceedings of The First International Workshop on Service Assurance with Partial and Intermittent Resources (SAPIR 2004)*, August 2004.

[10] S. Merugu, M. Ammar, and E. Zegura, “Space-time routing in wireless networks with predictable mobility,” Technical Report GIT-CC-04-07, College of computing, Georgia Institute of Technology, 2004.

[11] W. Navidi and T. Camp, “Stationary distributions for the random waypoint model,” *IEEE Transactions on Mobile Computing*, vol. 3, no. 1, pp. 99–108, 2004.

[12] A. Pentland, R. Fletcher, and A. Hasson, “A road to universal broadband connectivity,” In *Proceedings of the 2nd International Conference on Open Collaborative Design for Sustainable Innovation (dyd 02)*, December 2002.

[13] C.-C. Shen, G. Borkar, S. Rajagopalan, and C. Jaikaeo, “Interrogation-based relay routing for ad hoc satellite networks,” In *Proceedings of IEEE Globecom 2002*, November 2002.

[14] A. Vahdat and D. Becker, “Epidemic routing for partially connected ad hoc networks,” Technical Report CS-200006, Duke University, April 2000.

[15] K.-C. Wang and P. Ramananthan, “Collaborative sensing using sensors of uncoordinated mobility,” In *To appear in Proceedings of International Conference on Distributed Computing in Sensor Systems*, 2005.

[16] W. Zhao, M. Ammar, and E. Zegura, “A message ferrying approach for data delivery in sparse mobile ad hoc networks,” In *Proceedings of MobiHoc*, May 2004.

[17] W. Zhao, M. Ammar, and E. Zegura, “Controlling the mobility of multiple data transport ferries in a delay tolerant network,” In *To appear in IEEE INFOCOM 2005*, March 2005.