

Message Prioritization of Epidemic Forwarding in Delay-Tolerant Networks

Xu Liu, Yuanzhu Chen, Cheng Li

Wireless Networking and Mobile Computing Laboratory
Memorial University of Newfoundland
Canada

Walter Taylor and Jason H. Moore
Computational Genetics Laboratory
Giesel School of Medicine
Dartmouth College
USA

Abstract—Delay-tolerant networking is the innovation of extending the capabilities of communication networks in the temporal domain. It is an essential technology to enhance data communications in extreme scenarios, where crucial messages must be transferred with the infrastructure rendered unavailable or hostile. In this work, we consider the problem of disseminating a large number of messages in such networks. With the sparse and intermittently connected topology and with the unreliable and low-rate radios, the strategy of which messages to transfer first and in what order is a determinant of performance here. We compare a few such message prioritization methods using computer simulation and observe their performance in terms how widely and quickly information can be distributed across the network.

I. INTRODUCTION

Data communication networks connect computing devices with wired or wireless links to fulfill automated or human-initiated information exchange. For such networks to scale as the number of devices in it increases, we allow messages to traverse multiple communication links to travel from its source node to the destination. Such a “store and forward” technique is the essence of how the Internet is able to support billions of computers and users. This effectively extends the scope of communication networks spatially. In contrast, recent research on Delay-Tolerant Networking (DTN) [1], [13] has been exploring the possibility of extending communication networks temporally. With the advancement at the wireless networking and mobile computing frontier, mobile devices can be used to “store, carry, and forward” data when they roam around. That is, even without cellular infrastructure and only relying on short-range radios, e.g. Wi-Fi and Bluetooth, a set of sparsely deployed mobile nodes or handheld devices can be used to transfer data automatically if such data are not time-sensitive. The DTN technology can be useful in many scenarios, such as mobile sensor networks, disaster recovery, social networking, etc., to name a few.

At the center of DTN is the requirement that intermediate nodes should buffer messages for an extended period of time. This is in stark contrast to IP [14], where nodes simply discards messages when they are found undeliverable. To achieve a high or even acceptable deliver ratio, there is a significant burden on buffer space. Furthermore, because of the transient, unpredictable, low-bandwidth, and unreliable nature of device contact, efficacious utilization of such opportunities for nodes to exchange data is crucial. Therefore, tradeoff and

optimization in resource management is the major challenge in operating a DTN [5].

In the seminal work of Epidemic Routing [19], when two nodes (or devices) come into transmission range of each other, they conduct a 3-way handshake for one device to send messages to the other. In particular, when nodes A and B discover each other, node A sends an array of the IDs of all of its known messages to B . After receive this array and comparing it to its own set of buffered messages, node B replies with a subset of these message IDs as request indicating that these are the messages that A has but B does not. In the third message, node A sends the requested message to B . This process is triggered every time two nodes come close to each other. Apparently, epidemic routing can be a foundation for fulfilling both unicast and broadcast data transfer services.

In this article, we study the particular operation of information dissemination (i.e. message broadcasting) in DTNs. We do not assume any temporal self-similarity in the mobility of devices, so we do not rely on extrapolating previous contact information. We focus on the case where the number of buffered messages is so large that any form of compact representation of all stored messages, such as digest or ID, would not be accommodated in a single handshake packet. Therefore, when a node advertises about the messages it has received so far, it must pick strategically a subset of them to fit in a single packet. Such a strategy affects how quickly messages can be disseminated to a large number of nodes. We implemented the message dissemination framework with three prioritization approaches in The ONE [4], and tested them with one baseline approach. We observe that a well designed message prioritization method can significantly expedite such a broadcast service in DTNs.

The rest of this article is organized as follows. We review some of the most salient related research in Section II. In Section III, we discuss three message prioritization approach and one baseline method. We report experiment settings and results about these prioritization methods in Section IV and conclude the article with future research in Section V.

II. RELATED WORK

Research on DTN started from the Interplanetary Networking project at JPL [3]. The networking problem in such a scenario considers predictable mobility of space probes and surface stations, where the feedback loop can take a very long

time to complete due to both signal propagation delay and obstacles of other celestial bodies. In a more general setting, the mobility of communication devices is unpredictable, so scheduling networking activities in a deterministic fashion is no longer feasible. A great deal of research has been done on data transfer in such a framework to fulfill the simple goal of moving data from the source to its destination. A number of excellent reviews and vision articles have been published on the architecture and protocol aspects of delay-tolerant networks [11], [13], [20], [21].

The two most important, and yet distinct operations at the Network Layer are data forwarding and routing [14]. Forwarding regulates how packets are taken from one link and put on another. Routing determines which path a data packet should follow from the source node to the destination. The latter essentially provides the former with control input. Here, we stick to the term of data forwarding although it is also sometimes referred to as routing in literature.

Data forwarding in unpredictable DTN is more or less inspired by Epidemic Routing [19]. There, the authors are interested in transferring messages to their destinations as quickly as they can at the cost of using a large amount of network resources consumed by making many copies of the messages. Subsequent work on unicast data, where a message has a sole destination, explores the tradeoff between the data transfer performance, in terms of latency and delivery ratio, and resource consumption. For example, Spray and Wait [18] regulates the number of copies a message using a single control parameter.

When assuming that historical contact information would imply a similar pattern in future, nodes can utilize such observation to construct some sort of utility function to decide which node in its proximity might help forwarding its messages more effectively. This approach was initially explored in PROPHET [15] and MaxProp [6] most noticeably. When historical contact is further distilled with social network analysis methods, nodes can make more sophisticated forwarding decisions taking more factors into consideration. Such an approach is exemplified in BUBBLE-Rap [10], Delegation Forwarding [9], SimBet [8], and CAR [17].

Data forwarding techniques aside, researchers in data communications have also been on a quest for the killer applications of this groundbreaking technology for years [16]. Although such a quest is far from satisfactory, there have been a number of interesting applications in infrastructureless computer networking, such as IPN [3], Huggle [2], ZebraNet [12] and iSNAC [7], to name a few. Among these, iSNAC is a mobile social networking iPad application that focuses on broadcasting messages to help conference attendees to share information effectively.

III. MESSAGE PRIORITIZATION

In this work, we are interested in the problem of disseminating messages from a particular source to all other nodes in the network. We consider an intermittently connected network, where a set of sparsely deployed nodes roam around without assuming any predictability. Each node periodically injects a message into the network intended to all other nodes. We take a similar approach to Epidemic Routing in that, when two nodes

come into range of each other, one node can transfer a set of messages to the other via a 3-way handshake sequence. The difference, however, is that we must pick and choose which messages to include in a short advertisement packet so that the system has a high overall throughput. Specifically, when node A discovers node B , it sends a *summary vector* SV_A . In the original Epidemic Routing, SV_A contains the IDs of all the message that A stores in its buffer. In our solution, it is a subset of these messages because there can be many of them after the system has been up running for some time. After receiving SV_A , node B replies with $SV_A - M_B$, where M_B is the set of all messages stored at node B . As such, node B essentially tells A which messages from A would potentially enrich B 's collection. Next, node A retrieves messages in $SV_A - M_B$ from its storage and sends them to node B in a burst to complete the handshake. If the two nodes are still within range τ seconds after the handshake, they will start another round of handshake to transfer more messages.

Apparently, given that nodes typically store more messages than that can fit in a single handshake packet, the strategy taken to include which messages in the advertisement and in what order affects the system performance significantly. We call such a strategy *message prioritization*. In this work, we are interested in a few simple, and yet very different such methods. In all methods, we assume that node A fills the advertisement packet with l digests of some of its stored messages.

- 1) Round robin — Node A maintains a FIFO queue of the messages it has received and generated so far, i.e. by the time it is injected into the network. It circulates through the queue to compile the message digests using a pointer. When it is about to initiate a handshake, it processes l messages and advance the pointer accordingly. Here, the node maintains separate pointers for different nodes. Note that as the system continues, the time it takes to finish a round becomes longer, and when it does, it starts from the head of the queue again.
- 2) Tiered — Messages stored at a node are ranked according to three quantities to favor new, short messages, i.e. forward history, age, and length. The fewer times it has been forwarded till reaching this node A , the later it was created, and the shorter it is in length, it is ranked higher in the storage queue. These ranked messages are split into three segments of equal number of messages, the upper, middle, and lower tiers. Three separate round-robin schedules are executed on the tiers. The upper tier has three opportunities to send an advertisement containing l digests of its own, the middle tier has two, and the lower tier has one. As such, the system helps newly injected message to spread in the network more quickly.
- 3) Oblivious — Node A maintains a FIFO queue of all messages by the time they are injected into the network as in Round robin. When the node needs to create an advertisement packet, it simply takes the last l messages in the queue. In this method, the node never looks back after it has past a message in the queue, thus, always rigidly favoring the latest messages in the system.

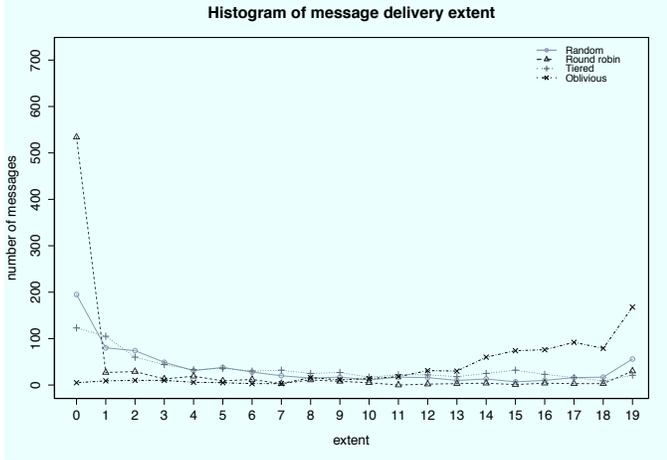


Fig. 3. Message delivery extent in 20-node network

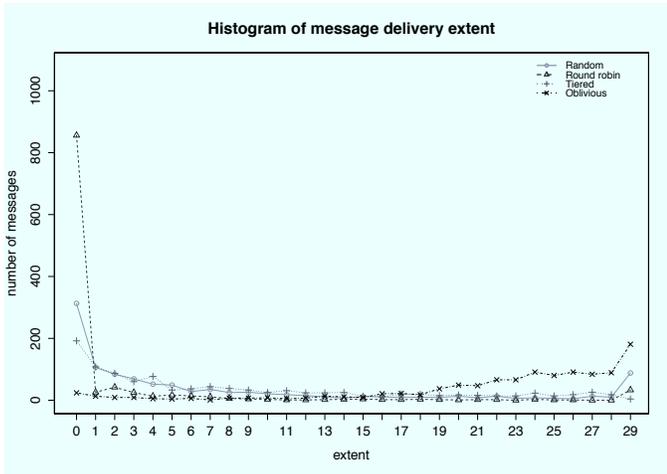


Fig. 4. Message delivery extent in 30-node network

Next, we turn our attention to how fast messages can be broadcast in the network using a generalized notion of latency, called *progression*. For a given message m in an n -node network, we use the vector $\langle m_1, m_2, \dots, m_{n-1} \rangle$ to denote the time it took to reach the i th other node ($i = 1, 2, \dots, n-1$). For the simulation of each of the message prioritization methods in a 10-node network, we summarize the message progression in a separate plot in the top half of Figure 6. Statistics shown in these plots include median, 25/75-quantile, 95% confidence, and outliers. In the bottom plot of the figure, we have the medians of the four methods together. Figures 7 and 8 present the same information for simulation in 20 and 30-node networks. We observe that the message progression rate of Oblivious is about an order of magnitude faster than the other methods, indicating that it is very effective directing messages.

V. CONCLUDING REMARKS

We are interested in message dissemination in delay-tolerant networks consisting of mobile devices roaming in a relatively confined area, such as a university building or conference facility. We had demonstrated some preliminary results previously in [7], where we deployed 12 iPads to

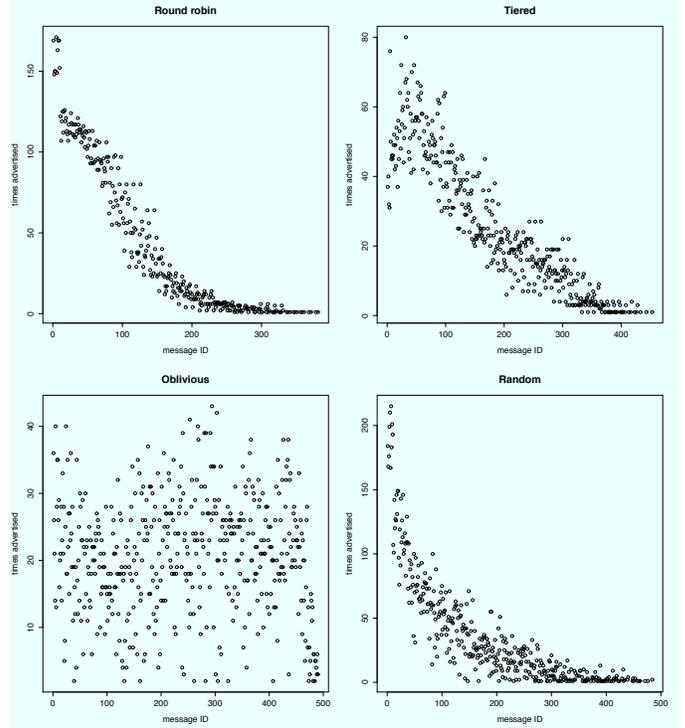


Fig. 5. Number of times a message is advertised in 10-node network

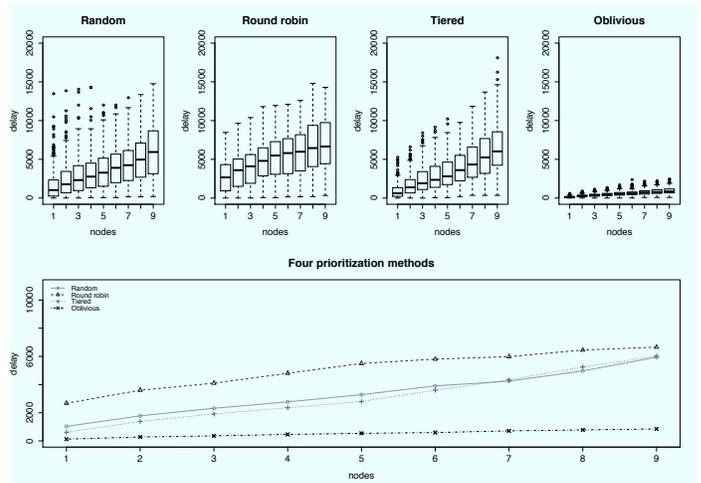


Fig. 6. Message delivery progression in 10-node network

investigate an infrastructureless social networking at the 61st Annual Meeting of the American Society of Human Genetics. In this article, we reported computer simulation results on a variety of message dissemination methods for a comparative study of their relative performance. With emphasis on newly injected messages, the Oblivious method is able to effectively distribute message across the network quickly. Our next step is to port these methods to actual mobile devices and test them at the same, real venue. By comparing the results to those reported here, we will be able to fine-tune some parameters in The ONE [4], so that we can use the simulator to test larger networks with better confidence.

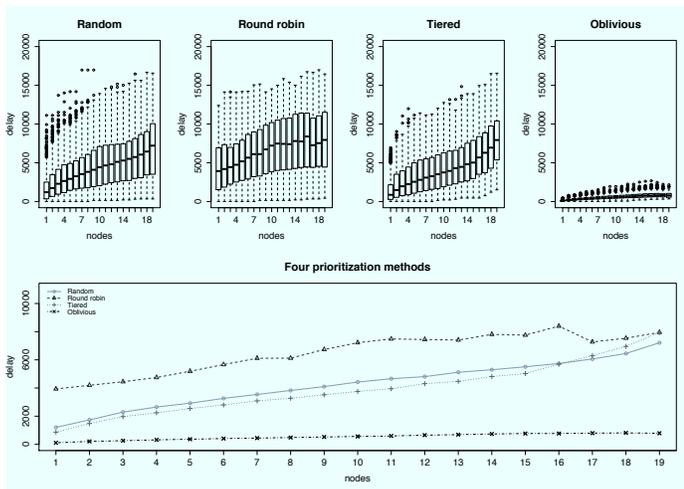


Fig. 7. Message delivery progression in 20-node network

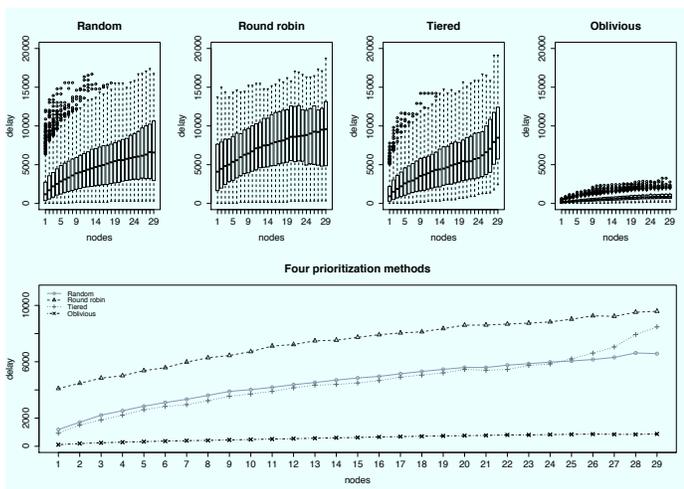


Fig. 8. Message delivery progression in 30-node network

ACKNOWLEDGEMENT

This work was supported by the Natural Sciences and Engineering Research Council (NSERC, Canada) Discovery Grant 327667-2010, and by National Institutes of Health (USA) grants R01-LM009012, R01-LM010098, and R01-AI59694.

REFERENCES

- [1] Delay tolerant networking research group. <http://www.dtnrg.org/wiki/Home>.
- [2] Haggle project. <http://code.google.com/p/haggle/>.
- [3] The interplanetary network (IPN). <http://tmo.jpl.nasa.gov/>.
- [4] The opportunistic network environment (ONE) simulator. <http://www.netlab.tkk.fi/tutkimus/dtn/theone/>.
- [5] Aruna Balasubramanian, Brian Levine, and Arun Venkataramani. DTN routing as a resource allocation problem. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 373–384, New York, NY, USA, 8 2007. ACM.
- [6] John Burgess, Brian Gallagher, David Jensen, and Brian Neil Levine. MaxProp: Routing for vehicle-based disruption-tolerant networks. In *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1–11, 2006.

- [7] Yuanchu Chen, Walter Taylor, Sam Coxon, and Jason H. Moore. isnac: Infrastructureless social networking at academic conferences. In *Demo at the 31st IEEE International Conference on Computer Communications (INFOCOM)*. IEEE, 2012.
- [8] Elizabeth M. Daly and Mads Haahr. Social network analysis for routing in disconnected delay-tolerant networks. In *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 32–40. ACM, 2007.
- [9] Vijay Erramilli, Mark Corvella, Augustin Chaintreau, and Christophe Diot. Delegation forwarding. In *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 251–260. ACM, 2008.
- [10] Pan Hui, Jon Crowcroft, and Eiko Yoneki. BUBBLE Rap: Social-based forwarding in delay tolerant networks. In *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 241–250. ACM, 2008.
- [11] Sushant Jain, Kevin Fall, and Rabin Patra. Routing in a delay tolerant network. In *Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 145–158. ACM, 2004.
- [12] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li Shiuian Peh, and Daniel Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with ZebraNet. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS X)*, pages 96–107, 2002.
- [13] Maurice J. Khabbaz, Chadi M. Assi, and Wissam F. Fawaz. Disruption-tolerant networking: A comprehensive survey on recent developments and persisting challenges. *IEEE Communications Surveys & Tutorials*, 14(2):607–640, 2012.
- [14] James F. Kurose and Keith W. Ross. *Computer Networking: A Top-Down Approach*. Pearson, Upper Saddle River, NJ, USA, 6th edition, 2012.
- [15] Anders Lindgren, Avri Doria, and Olov Schelén. Probabilistic routing in intermittently connected networks. *SIGMOBILE Mobile Computing and Communications Review*, 7(3):19–20, July 2003.
- [16] Anders Lindgren and Pan Hui. The quest for a killer app for opportunistic and delay tolerant networks. In *Proceedings of the 4th ACM Workshop on Challenged Networks (CHANTS)*, pages 59–66. ACM, 2009.
- [17] Mirco Musolesi and Cecilia Mascolo. CAR: Context-aware adaptive routing for delay-tolerant mobile networks. *IEEE Transactions on Mobile Computing*, 8(2):246–260, February 2009.
- [18] Thrasyvoulos Spyropoulos, Konstantinos Psounis, and Cauligi S. Raghavendra. Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In *Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, pages 252–259. ACM, 2005.
- [19] Amin Vahdat and David Becker David. Epidemic routing for partially-connected ad hoc networks. Technical Report CS-200006, Duke University, 2000.
- [20] Zhensheng Zhang. Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: overview and challenges. *IEEE Communications Surveys & Tutorials*, 8(1):24–37, 2006.
- [21] Zhensheng Zhang and Qian Zhang. Delay/disruption tolerant mobile ad hoc networks: latest developments. *Wireless Communications and Mobile Computing*, 7(10):1219–1232, December 2007.