

Encounter-Based Routing in DTNs

Samuel C. Nelson, Mehedi Bakht, Robin Kravets^a Albert F Harris III^b
{snelso20,mbakht2,rhk}@cs.uiuc.edu afh@cresis.ku.edu

^aDepartment of Computer Science, University of Illinois at Urbana-Champaign

^bCenter for Remote Measurement of Ice Sheets, University of Kansas

(Problem and Motivation)

I. Introduction

Delay and disruption tolerant networks (DTNs) transport application data by creating a “store and forward” network where no infrastructure exists. Although end-to-end connectivity may not be available between two nodes, DTN routing protocols take advantage of temporal paths created in the network as nodes encounter their neighbors and exchange messages they have been asked to forward. Since there are no guarantees that a route will ever be available, many current DTN routing protocols apply epidemic-style techniques [9], leveraging the fact that an increased number of copies of a particular message in the network should improve the probability that the message will reach its intended destination. However, such techniques come at a high price in terms of network resources, resulting in the rapid depletion of buffer space and energy on resource-limited devices, the rapid depletion of available bandwidth, and the potential to greatly increase end-to-end delay.

A number of routing protocols have been proposed to enable data delivery in such challenging environments [2, 3, 4, 5, 7, 8]. However, many of these protocols trade overhead and computational complexity for increased successful delivery. This overhead expresses itself as more traffic in the network creating more contention in clusters of high connectivity and increased energy consumption for nodes exchanging messages.

One method to mitigate this effect is to identify properties in the network that could be used to make intelligent forwarding and message replication decisions. In environments targeted by DTNs, such as disaster scenarios and certain vehicular networks, different classes of nodes naturally tend to have more node encounters than others. The main contribution of our research takes this network property and uses it to design a DTN routing protocol that uses only local observations about a node's environment. Our protocol, Encounter-Based Routing (EBR), uses an encounter-based metric for optimization of message passing that maximizes message delivery while minimizing overhead both in terms of extra traffic injected into the network and control overhead, as well as minimizing latency. We evaluate EBR using multiple mobility models and show a high packet delivery ratio while incurring extremely low overhead.

(Background and Related Work)

II. DTN Routing Protocol Taxonomy

DTN routing protocols can be classified into two high-level approaches [2]. *Forwarding-based* protocols only keep one copy of a message in the network and attempt to forward that copy toward the destination at each encounter. In contrast, *replication-based* protocols insert multiple copies, or replicas, of a message into the network to increase the probability of message delivery. Essentially, replication-based protocols leverage a trade-off between resource usage (*e.g.*, node memory and bandwidth) and probability of message delivery. Replication-based protocols can be further separated into two classes based on the number of replicas created. *Flooding-based* protocols attempt to send a replica of each message to as many nodes as possible (*e.g.*, MaxProp [3], RAPID [2], and Prophet [5]), whereas *quota-based* protocols intentionally limit the number of replicas (*e.g.*, Spray and Wait [7] and Spray and Focus [8]).

Forwarding-based protocols are currently limited in their effectiveness due to the difficulty or impossibility of finding a path from the source to the destination with only one copy, and hence replication-based protocols are extremely popular. One main problem with flooding-based replication protocols is their high demand on network resources, such as storage and bandwidth. This fact led to some work in developing quota-based protocols, which are much better stewards of network resources than their flooding-based counterparts. However, one possible criticism is their inability to successfully deliver a comparable amount of messages, in certain cases. We show

this to be false by developing a quota-based protocol using an encounter-based routing metric that has extremely low routing overhead, while maintaining delivery ratios better than or comparable to current flooding-based protocols.

(Approach / Uniqueness of Approach)

III. Encounter-based Routing (EBR)

We present Encounter-based Routing (EBR), which is a quota-based DTN routing protocol that achieves a high delivery ratio comparable to flooding-based protocols, while maintaining extremely low network overhead. This improvement in delivery ratio is accomplished by taking advantage of the following observed mobility property of certain networks: *the future rate of node encounters can be roughly predicted by past data*. This property is useful because nodes that experience a large number of encounters are more likely to successfully pass the message along to the final destination than nodes that only infrequently encounter others. Many networks experience this phenomenon; examples include disaster recovery networks, where ambulances and police tend to be more

mobile and bridge more cluster gaps than civilians [6], and vehicular-based networks, where certain vehicles take popular routes. Utilizing this observation, EBR bases routing decisions on a measure of a node's rate of encounters, showing preference to message exchanges with nodes that have high encounter rates.

In EBR, information about a node's rate of encounter is a purely local metric and can be tracked using a small number of variables. Therefore, EBR is able to maintain very low state overhead, as compared to other protocols that can require up to $O(n)$ routing messages exchanged during *every* contact connection, and $O(n^2)$ routing state locally stored (e.g., MaxProp [3], Prophet [5]). A further strength of EBR is that its message replication rules are simple to understand and implement, as opposed to complex rules found in many protocols, minimizing the chance of bugs and reducing computational complexity.

III.A. Algorithm

Every node running the EBR protocol is responsible for maintaining its past rate of encounter average, which is used to predict future encounter rates. When two nodes meet, the relative ratio of their respective rates of encounter determines the appropriate fraction of message replicas the nodes should exchange. The primary purpose of tracking the rate of encounter is to intelligently decide how many replicas of a message a node should transfer during a contact opportunity.

To track a node's rate of encounter, it maintains two pieces of local information: an encounter value (EV), and a current window counter (CWC). EV represents the node's past rate of encounters as an exponentially weighted moving average, while CWC is used to obtain information about the number of encounters in the current time interval. EV is periodically updated to account for the most recent CWC in which rate of encounter information was obtained. Updates to EV are computed as follows:



This exponentially weighted moving average places an emphasis proportional to α on the most recent complete CWC. Updating CWC is straightforward: for every encounter, CWC is incremented. When the current window update interval has expired, EV is updated and the CWC is reset to zero. In our experiments, we found an α of 0.85 and update interval of around 30 seconds allow for reasonable results.

Since EV represents a prediction of the future rate of encounters for each node per time interval, the node with the highest EV represents a higher probability of successful message delivery. Therefore, when two nodes meet they compare their EVs. The number of replicas of a message transferred during a contact opportunity is proportional to the ratio of the EVs of the nodes. For two nodes A and B, where $EV_A < EV_B$, for every message M_i , node A sends



replicas of M_i , where m_i is the total number of M_i replicas stored at node A. For example, assume node A has 4 replicas of a message M_1 and 8 replicas of a message M_2 . Furthermore, assume node A, with $EV_A = 5$, comes in contact with node B, with $EV_B = 15$. Since $EV_A < EV_B$, node A sends $\frac{1}{3}$ of the replicas of each message. Therefore, node A transmits 3 replicas of M_1 and 6 replicas of M_2 .

Algorithm 1 presents the basic form of the EBR protocol, where W_i represent the current window update interval parameter.



III.B. Generalizing EBR

In this section, we prove that EBR adheres to the definition of a quota-based protocol (as described in Section II) and show the relevant bounds, both for the simple version, where L , the maximum number of replicas of a message, is discrete, and for a more general version, allowing the use of probabilistic L values.

For discrete L values, it is easy to show that EBR is quota-based. Along with its data, every message contains a value indicating the maximum number of replicas into which this current message is allowed to be split. As an example, assume an application at node A creates a message with the maximum allowable replicas set to 10. Assume node A encounters node B and, based on the EBR protocol described in Section III.A, wishes to transmit 8 replicas. Then, A creates a copy of the message for node B and assigns B 's maximum allowable replicas to 8. Furthermore, A resets its maximum allowable replicas to 2. Continuing this procedure in a recursive fashion maintains the bound set by the initial message.

However, L values are not limited to a discrete maximum number of replicas. The discrete structure can easily be relaxed into a probabilistic structure, while maintaining meaningful (yet probabilistic) bounds. Probabilistic L values can allow for less sensitivity to exact network conditions. When using discrete L values, changes to the initial number of message replicas allows for a fundamental tradeoff between message delivery ratio, goodput, and average latency (see Section IV). Using probabilistic L values and changing the variance and mean can allow applications to compromise and not require exact decisions about the number of allowable replicas.

While any distribution may be used in this probabilistic model, the Gaussian distribution allows for immediate, eloquent properties that help establish the bound on the number of messages in the network. In this case, the application specifies the mean and variance of the distribution, instead of a discrete number. Assume a node A wishes to split the message M into two replicas, M_A and M_B . Node A must follow the following EBR message splitting rule:

If $M \sim N(\mu, \sigma^2)$, then it can only be split into $M_A \sim N(\mu_A, \sigma_A^2)$ and $M_B \sim N(\mu_B, \sigma_B^2)$ such that $\mu = \mu_A + \mu_B$ and $\sigma^2 = \sigma_A^2 + \sigma_B^2$.

For example, a message with mean 10 and variance 5 may be split into two messages, one with mean 8 and variance 4, and one with mean 2 and variance 1. It may not, however, be split into a message of mean 8 and variance 4, and one with mean 7 and variance 1. As a further note, EBR maintains the ratio of mean to variance for all message splits.

This message splitting rule preserves the Gaussian distribution for the two newly created replicas. This is due to a result from statistics known as Cramer's Theorem:



We now demonstrate that this general version of EBR is a quota-based replication protocol, and establish an upper bound, by proving the following theorem:

Theorem III.1 *Let S be a schedule of future message creations. Let t be an arbitrary future time. Assume M_1, M_2, \dots, M_i in S are all the messages created before time t . Assume each message M_i has a Gaussian random variable (for notational ease, we refer to this directly as the message M_i), with mean μ_i and variance σ_i^2 , that represents the maximum number of replicas the current message is allowed to be split into.*

The upper bound on the maximum number of message replicas in the system is:



Proof: Let U be the sum of all message replicas in the system. Assuming messages never split, there will be i messages in the system, each with mean μ_i and variance σ_i^2 . We utilize the following rule of linearity for Gaussian distributions (the converse of Cramer's Theorem):



Now assume a message, $M_j \sim N(\mu_j, \sigma_j^2)$ is split into $M_{j1} \sim N(\mu_{j1}, \sigma_{j1}^2)$ and $M_{j2} \sim N(\mu_{j2}, \sigma_{j2}^2)$ such that $\mu_j = \mu_{j1} + \mu_{j2}$ and $\sigma_j^2 = \sigma_{j1}^2 + \sigma_{j2}^2$ (the message splitting rule of EBR). Then by the same linearity rules, $M_j = M_{j1} + M_{j2}$, leaving U unchanged. *QED*

One minor issue to address is that the statistical rules and theorems each assume true Gaussian distributions. However, it does not make sense in our system for a message M to hold a negative value. The probability of this occurring can be made sufficiently small by forcing the application to choose sufficiently low variances for corresponding means (which can never be below zero).

(Results and Contributions)

IV. Evaluation

The primary goal of our evaluation is to show that EBR achieves a high message delivery ratio and good latency, while maintaining extremely low overhead. To perform our evaluation, we use the Opportunistic Network Environment simulator (ONE) [1], a simulation environment designed specifically for DTNs. We evaluate EBR against five other popular protocols: (1) basic epidemic [9], (2) Prophet [5], (3) Spray and Wait [7], (4) Spray and Focus [8], and (5) MaxProp [3].

To evaluate the protocols in a realistic vehicular-based DTN, we utilize the map-driven mobility model of ONE to limit node movement to actual streets found on an imported map, an approximate 5 km x 3 km section of downtown Helsinki, Finland. Approximately 15% of the nodes are configured to follow pre-defined routes with speed between 7 and 10 m/s. The rest of the nodes are divided into four groups and assigned different probabilities of picking the next destination from a particular set of points to simulate the phenomenon that people often visit certain areas of a city more frequently than others based on their profession, age and other factors. The speed of these nodes varies between 2.7 and 13.9 m/s.

We also simulate the routing protocols with a traditional random waypoint model in an area of size 3 km by 3 km. Node speed is varied between 0.5 and 1.5 m/s and the pause time at destination is distributed between 0 and 120 seconds.

The transmission range of each node is 250 m. We vary the number of nodes in the network starting at 26, followed by 51 to 251 in increments of 50. The packet size and buffer space are kept constant at 25 KB and 1 MB respectively. Each simulation lasts for one simulated hour. Except for a small number of MaxProp data points, each point is the average of at least 10 runs, with 95% confidence intervals displayed.

The vehicular mobility model exhibits the property that past information on rate-of-encounters is a good estimator for future rate-of-encounters. Since it fits perfectly into the assumptions of EBR, EBR performs extremely well in this model, especially in terms of message delivery ratio (MDR) (see Figure 1(a)). EBR is also, by far, the most resource friendly, as shown by the goodput metric (see Figure 1(c)). While EBR seems to have unfavorable delay, this is, in part, due to a high MDR (see Figure 1(b)). Since delay is computed only over messages that have been delivered, it is deceptive to view delay alone since many protocols quickly deliver messages that take a small number of hops, and do not deliver most high-hop messages.

In contrast, the random waypoint model lacks heterogeneity in node mobility, a property that EBR was designed to leverage. As a result, EBR does not perform as well in this model. However it still beats other protocols clearly in terms of MDR. The goodput metric also strongly favors EBR (see Figure 2(c)). Notice that in both the vehicular mobility model and the random waypoint mobility model, EBR substantially beats the other quota-based protocols (specifically Spray and Wait), and obtains more than twice their goodput in many cases. This is in addition to consistently performing best in terms of message delivery ratio.



References

- [1] Opportunistic network environment ONE. <http://www.netlab.tkk.fi/tutkimus/dtn/theone/>.

- [2] Aruna Balasubramanian, Brian Neil Levine, and Arun Venkataramani. DTN routing as a resource allocation problem. In *Proc. ACM SIGCOMM*, August 2007.
- [3] John Burgess, Brian Gallagher, David Jensen, and Brian Neil Levine. MaxProp: Routing for vehicle-based disruption-tolerant networks. In *Proc. IEEE INFOCOM*, April 2006.
- [4] Vijay Erramilli, Augustin Chaintreau, Mark Crovella, and Christophe Diot. Diversity of forwarding paths in pocket switched networks. In *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 161–174, New York, NY, USA, 2007. ACM.
- [5] A. Lindgren, A. Doria, and O. Scheln. Probabilistic routing in intermittently connected networks. In *Proceedings of the Fourth ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2003)*, 2003.
- [6] Samuel C. Nelson, Albert F. Harris, and Robin Kravets. Event-driven, role-based mobility in disaster recovery networks. In *CHANTS 07: Proceedings of the second workshop on Challenged networks*, 2007.
- [7] Thrasylvoulos Spyropoulos, Konstantinos Psounis, and Cauligi S. Raghavendra. Spray and wait: An efficient routing scheme for intermittently connected mobile networks. In *WDTN '05: Proceeding of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, 2005.
- [8] Thrasylvoulos Spyropoulos, Konstantinos Psounis, and Cauligi S. Raghavendra. Spray and focus: Efficient mobility-assisted routing for heterogeneous and correlated mobility. In *Fifth Annual IEEE International Conference on Pervasive Computing and Communications Workshops*, 2007.
- [9] Amin Vahdat and David Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-2000-06, Department of Computer Science, Duke University, apr 2000.