

Replication Routing in DTNs: A Resource Allocation Approach

Aruna Balasubramanian Brian Neil Levine Arun Venkataramani
 Department of Computer Science, University of Massachusetts, Amherst, USA 01003
 {arunab, brian, arun}@cs.umass.edu

Abstract—Routing protocols for disruption-tolerant networks (DTNs) use a variety of mechanisms, including discovering the meeting probabilities among nodes, packet replication, and network coding. The primary focus of these mechanisms is to increase the likelihood of finding a path with limited information, and so these approaches have only an *incidental* effect on such routing metrics as maximum or average delivery delay. In this paper, we present RAPID, an *intentional* DTN routing protocol that can optimize a specific routing metric such as the worst-case delivery delay or the fraction of packets that are delivered within a deadline. The key insight is to treat DTN routing as a resource allocation problem that translates the routing metric into per-packet utilities which determine how packets should be replicated in the system. We evaluate RAPID rigorously through a prototype deployed over a vehicular DTN testbed of 40 buses and simulations based on real traces. To our knowledge, this is the first paper to report on a routing protocol deployed on a real outdoor DTN. Our results suggest that RAPID significantly outperforms existing routing protocols for several metrics. We also show empirically that for small loads, RAPID is within 10% of the optimal performance.

I. INTRODUCTION

Disruption-tolerant networks (DTNs) enable transfer of data when mobile nodes are connected only intermittently. Applications of DTNs include large-scale disaster recovery networks, sensor networks for ecological monitoring [34], ocean sensor networks [26], [22], vehicular networks [24], [7], and projects such as TIER [2], Digital Study Hall [14], and One Laptop Per Child [1] to benefit developing nations. Intermittent connectivity can be a result of mobility, power management, wireless range, sparsity, or malicious attacks. The inherent uncertainty about network conditions make routing in DTNs a challenging problem.

The primary focus of many existing DTN routing protocols is to increase the likelihood of finding a path with extremely limited information. To discover such a path, a variety of mechanisms are used, including estimating node meeting probabilities, packet replication, network coding, placement of stationary waypoint stores, and leveraging prior knowledge of mobility patterns. Unfortunately, the burden of finding even one path is so great that existing approaches have only an *incidental* rather than an *intentional* effect on such routing *metrics* as worst-case delivery latency, average delay, or percentage of packets delivered. This disconnect between application needs and routing protocols hinders deployment of DTN applications. Currently, it is difficult to drive the routing layer of a DTN by specifying priorities, deadlines, or cost constraints. For

example, a simple news and information application is better served by maximizing the number of news stories delivered before they are outdated, rather than eventually delivering all stories.

In this paper, we formulate the DTN routing problem as a *resource allocation problem*. The protocol we describe, called RAPID (Resource Allocation Protocol for Intentional DTN) routing, allocates resources to packets to optimize an administrator-specified routing metric. At each transfer opportunity, a RAPID node replicates or allocates bandwidth resource to a set of packets in its buffer, in order to optimize the given routing metric. Packets are delivered through opportunistic replication, until a copy reaches the destination.

RAPID makes the allocation decision by first translating the routing metric to a per-packet utility. DTNs are resource-constrained networks in terms of transfer bandwidth, energy, and storage; allocating resources to replicas without careful attention to available resources can cause more harm than good. Therefore, a RAPID node replicates packets in the order of their marginal utility of replication, i.e., the first packet to be replicated is the one that provides the highest increase in utility per unit resource used. We show how RAPID can use this simple approach to optimize three different routing metrics: average delay, worst-case delay, and the number of packets delivered before a deadline.

RAPID loosely tracks network resources through a control plane to assimilate a local view of the global network state. To this end, RAPID uses an in-band *control channel* to exchange network state information among nodes using a fraction of the available bandwidth, and uses the additional information to significantly improve routing performance. RAPID’s control channel builds on insights from previous work. For example, Jain et al. [18] suggest that DTN routing protocols that use more knowledge of network conditions perform better, and Burgess et al. [7] show that flooding acknowledgments improves delivery rates by removing useless packets from the network.

We present hardness results to substantiate RAPID’s heuristic approach. We prove that online algorithms without complete future knowledge and with unlimited computational power, or computationally limited algorithms with complete future knowledge, can be arbitrarily far from optimal.

We have built and deployed RAPID on a vehicular DTN testbed, DieselNet [7], that consists of 40 buses covering a 150 square-mile area around Amherst, MA. We collected 58 days of performance traces of the RAPID deployment. To our knowledge, this is the first paper to report on a routing

protocol deployed on a real outdoor DTN. Similar testbeds have deployed only flooding as a method of packet propagation [34]. We also conduct a simulation-based evaluation using real traces to stress-test and compare various protocols. We show that the performance results from our trace-driven simulation is within 1% of the real measurements with 95% confidence. We use this simulator to compare RAPID to four existing routing protocols [21], [29], [7] and random routing. We also compare the protocols using synthetic mobility models.

We evaluate the performance of RAPID for three different routing metrics: average delay, worst-case delay, and the number of packets delivered before a deadline. All experiments include the cost of RAPID’s control channel. Our experiments using trace-driven and synthetic mobility scenarios show that RAPID significantly outperforms the four routing protocols. For example, in trace-driven experiments under moderate-to-high loads, RAPID outperforms the second-best protocol by about 20% for all three metrics, while also delivering 15% more packets for the first two metrics. With *a priori* mobility information and moderate-to-high loads, RAPID outperforms random replication by about 50% for high packet loads. We also compare RAPID to an optimal protocol and show empirically that RAPID performs within 10% of optimal for low loads.

II. RELATED WORK

a) Replication versus Forwarding: We classify related existing DTN routing protocols as those that replicate packets and those that forward only a single copy. *Epidemic* routing protocols replicate packets at transfer opportunities hoping to find a path to a destination. However, naive flooding wastes resources and can severely degrade performance. Proposed protocols attempt to limit replication or otherwise clear useless packets in various ways: (i) using historic meeting information [13], [8], [7], [21]; (ii) removing useless packets using acknowledgments of delivered data [7]; (iii) using probabilistic mobility information to infer delivery [28]; (iv) replicating packets with a small probability [33]; (v) using network coding [32] and coding with redundancy [17]; and (vi) bounding the number of replicas of a packet [29], [28], [23].

In contrast, *forwarding* routing protocols maintain at most one copy of a packet in the network [18], [19], [31]. Jain et al. [18] propose a forwarding algorithm to minimize the average delay of packet delivery using oracles with varying degrees of future knowledge. Our deployment experience suggests that, even for a scheduled bus service, implementing the simplest oracle is difficult; connection opportunities are affected by many factors in practice including weather, radio interference, and system failure. Furthermore, we present formal hardness and empirical results to quantify the impact of not having complete knowledge.

Jones et al. [19] propose a link-state protocol based on epidemic propagation to disseminate global knowledge, but use a single path to forward a packet. Shah et al. [27] and Spyropoulos et al. [31] present an analytical framework for the forwarding-only case assuming a grid-based mobility model. They subsequently extend the model and propose a replication-based protocol, *Spray and Wait* [29]. The consensus appears

to be [29] that replicating packets can improve performance (and security [6]) over just forwarding, but risk degrading performance when resources are limited.

b) Incidental versus Intentional: Our position is that most existing schemes only have an *incidental* effect on desired performance metrics, including commonly evaluated metrics such as average delay or delivery probability. Therefore, the effect of a routing decision on the performance of a given resource constrained network scenario is unclear. For example, several existing DTN routing algorithms [29], [28], [23], [7] route packets using the number of replicas as the heuristic, but the effect of replication varies with different routing metrics. *Spray and Wait* [29] routes to reduce delay metric, but it does not take into account bandwidth or storage constraints. In contrast, routing in RAPID is *intentional* with respect to a given performance metric. RAPID explicitly calculates the effect of replication on the routing metric while accounting for resource constraints.

c) Resource Constraints: RAPID also differs from most previous work in its assumptions regarding resource constraints, routing policy, and mobility patterns. Table I shows a taxonomy of many existing DTN routing protocols based on assumptions about *bandwidth* available during transfer opportunities and the *storage* carried by nodes; both are either *finite* or *unlimited*. For each work, we state in parentheses the mobility model used. RAPID is a replication-based algorithm that assumes constraints on both storage and bandwidth (P5) — the most challenging and most practical problem space.

P1 and P2 are important to examine for valuable insights that theoretical tractability yields but are impractical for real DTNs with limited resources. Many studies [21], [13], [8], [28] analyze the case where storage at nodes is limited, but bandwidth is unlimited (P3). However, we find this scenario to be uncommon. Bandwidth is likely to be constrained for most typical DTN scenarios. Specifically, in mobile and vehicular DTNs, transfer opportunities are typically short-lived [16], [7].

We were unable to find other protocols in P5 except *MaxProp* [7] that assume limited storage and bandwidth. However, it is unclear how to optimize a specific routing metric using *MaxProp*, so we categorize it as an incidental routing protocol. Our experiments indicate that RAPID outperforms *MaxProp* for each metric that we evaluate.

Some theoretical works [35], [30], [28], [5] derive closed-form expressions for average delay and number of replicas in the system as a function of the number of nodes and mobility patterns. Although these analyses contributed to important insights in the design of RAPID, their assumptions about mobility patterns or unlimited resources were, in our experience, too restrictive to be applicable to practical settings.

III. THE RAPID PROTOCOL

A. System model

We model a DTN as a set of mobile nodes. Two nodes transfer data packets to each other when within communication range. During a transfer, the sender replicates packets while retaining a copy. A node can deliver packets to a destination node directly or via intermediate nodes, but packets may not

Problem	Storage	Bandwidth	Routing	Previous work (and mobility)
P1	Unlimited	Unlimited	Replication	Epidemic [23], Spray and Wait [29]: Constraint in the form of channel contention (Grid-based synthetic)
P2	Unlimited	Unlimited	Forwarding	Modified Dijkstra's et al. [18] (simple graph), MobySpace [20] (Powerlaw)
P3	Finite	Unlimited	Replication	Davis et al. [13] (Simple partitioning synthetic), SWIM [28] (Exponential), MV [8] (Community-based synthetic), Prophet [21] (Community-based synthetic)
P4	Finite	Finite	Forwarding	Jones et al. [19] (AP traces), Jain et al. [18] (Synthetic DTN topology)
P5	Finite	Finite	Replication	This paper (Vehicular DTN traces, exponential, and power law meeting probabilities, testbed deployment), MaxProp [7] (Vehicular DTN traces)

TABLE I
A CLASSIFICATION OF SOME RELATED WORK INTO DTN ROUTING SCENARIOS

be fragmented. There is limited storage and transfer bandwidth available to nodes. Destination nodes are assumed to have sufficient capacity to store delivered packets, so only storage for in-transit data is limited. Node meetings are assumed to be short-lived. The nodes are assumed to have sufficient computational capabilities as well as enough resources to maintain state information.

Formally, a DTN consists of a node meeting schedule and a workload. The node meeting schedule is a directed multigraph $G = (V, E)$, where V and E represent the set of nodes and edges, respectively. Each directed edge e between two nodes represents a meeting between them, and it is annotated with a tuple (t_e, s_e) , where t is the time and s is the size of the transfer opportunity. The workload is a set of packets $P = \{(u_1, v_1, s_1, t_1), (u_2, v_2, s_2, t_2), \dots\}$, where the i th tuple represents the source, destination, size, and time of creation (at the source), respectively, of packet i . The goal of a DTN routing algorithm is to deliver all packets using a feasible schedule of packet transfers, where *feasible* means that the total size of packets transferred during each opportunity is less than the size of the opportunity, always respecting storage constraints.

In comparison to Jain et al.[18] who model link properties as continuous functions of time, our model assumes discrete short-lived transfers; this makes the problem analytically more tractable and characterizes many practical DTNs well.

B. RAPID design

RAPID models DTN routing as a utility-driven resource allocation problem. A packet is routed by replicating it until a copy reaches the destination. The key question is: given limited bandwidth, how should packets be replicated in the network so as to optimize a specified *routing metric*? RAPID derives a per-packet *utility function* from the routing metric. At a transfer opportunity, it replicates a packet that locally results in the highest increase in utility.

Consider a routing metric such as *minimize average delay of packets*, the running example used in this section. The corresponding utility U_i of packet i is the negative of the expected delay to deliver i , i.e., the time i has already spent in the system plus the additional expected delay before i is delivered. Let δU_i denote the increase in U_i by replicating i and s_i denote the size of i . Then, RAPID replicates the packet

$D(i)$	Packet i 's expected delay = $T(i) + A(i)$
$T(i)$	Time since creation of i
$a(i)$	Random variable that determines the remaining time to deliver i
$A(i)$	Expected remaining time = $E[a(i)]$
M_{XZ}	Random variable that determines inter-meeting time between nodes X and Z

TABLE II
LIST OF COMMONLY USED VARIABLES.

with the highest value of $\delta U_i/s_i$ among packets in its buffer; in other words, the packet with the highest marginal utility.

In general, U_i is defined as the expected contribution of i to the given routing metric. For example, the metric *minimize average delay* is measured by summing the delay of packets. Accordingly, the utility of a packet is its expected delay. Thus, RAPID is a heuristic based on locally optimizing marginal utility, i.e., the expected increase in utility per unit resource used.

Using the marginal utility heuristic has some desirable properties. The marginal utility of replicating a packet to a node is low when (i) the packet has many replicas, or (ii) the node is a poor choice with respect to the routing metric, or (iii) the resources used do not justify the benefit. For example, if nodes meet each other uniformly, then a packet i with 6 replicas has lower marginal utility of replication compared to a packet j with just 2 replicas. On the other hand, if the peer is unlikely to meet j 's destination for a long time, then i may take priority over j .

RAPID has three core components: a *selection* algorithm, an *inference* algorithm, and a *control channel*. The selection algorithm is used to determine *which* packets to replicate at a transfer opportunity given their utilities. The inference algorithm is used to estimate the *utility* of a packet given the routing metric. The control channel propagates the necessary metadata required by the inference algorithm.

C. The selection algorithm

The RAPID protocol executes when two nodes are within radio range and have discovered one another. The protocol is symmetric; without loss of generality, we describe how node X determines which packets to transfer to node Y (refer to the box marked PROTOCOL RAPID).

PROTOCOL RAPID(X, Y):

- 1) *Initialization*: Obtain metadata from Y about packets in its buffer as well as metadata it collected over past meetings (detailed in Section IV-B).
- 2) *Direct delivery*: Deliver packets destined to Y in decreasing order of creation times.
- 3) *Replication*: For each packet i in node X 's buffer
 - a) If i is already in Y 's buffer (as determined from the metadata), ignore i .
 - b) Estimate marginal utility, $\delta U_i/s_i$, of replicating i to Y .
 - c) Replicate packets in decreasing order of marginal utility.
- 4) *Termination*: End transfer when out of radio range or all packets replicated.

RAPID also adapts to storage restrictions for in-transit data. If a node exhausts all available storage, packets with the lowest utility are deleted first as they contribute least to overall performance. However, a source never deletes its own packet unless it receives an acknowledgment for the packet.

D. Inference algorithm

Next, we describe how PROTOCOL RAPID can support specific metrics using an algorithm to infer utilities. Table II defines the relevant variables.

1) **Metric 1: Minimizing average delay**: To minimize the average delay of packets in the network we define the utility of a packet as

$$U_i = -D(i) \quad (1)$$

since the packet's expected delay is its contribution to the performance metric. RAPID attempts to greedily replicate the packet whose replication reduces the delay by the most among all packets in its buffer.

2) **Metric 2: Minimizing missed deadlines**: To minimize the number of packets that miss their deadlines, the utility is defined as the probability that the packet will be delivered within its deadline:

$$U_i = \begin{cases} P(a(i) < L(i) - T(i)), & L(i) > T(i) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where $L(i)$ is the packet life-time. A packet that has missed its deadline can no longer improve performance and is thus assigned a value of 0. The marginal utility is the improvement in the probability that the packet will be delivered within its deadline.

3) **Metric 3: Minimizing maximum delay**: To minimize the maximum delay of packets in the network, we define the utility U_i as

$$U_i = \begin{cases} -D(i), & D(i) \geq D(j) \quad \forall j \in S \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where S denotes the set of all packets in X 's buffer. Thus, U_i is the negative expected delay if i is a packet with the maximum expected delay among all packets held by Y . So, replication is useful only for the packet whose delay is maximum. For

the routing algorithm to be work conserving, RAPID computes utility for the packet whose delay is currently the maximum; i.e., once a packet with maximum delay is evaluated for replication, the utility of the remaining packets is recalculated using Eq. 3.

IV. ESTIMATING DELIVERY DELAY

How does a RAPID node estimate expected delay in Eqs. 1 and 3, or the probability of packet delivery within a deadline in Eq. 2? The expected delivery delay is the minimum expected time until any node with the replica of the packet delivers the packet; so a node needs to know which other nodes possess replicas of the packet and when they expect to meet the destination.

To estimate expected delay we assume that each node with the copy of the packet delivers the packet directly to the destination, ignoring the effect of further replications. This assumption simplifies the expected delay estimation, and we make this assumption only for networks with dense node meetings, where every node meets every other node. In Section IV-A2, we describe a modification to this assumption for networks with sparse node meetings. Estimating expected delay is nontrivial even with an accurate global snapshot of system state. For ease of exposition, we first present RAPID's estimation algorithm as if we had knowledge of the global system state, and then we present a practical distributed implementation.

A. Algorithm Estimate_Delay

A RAPID node uses the algorithm ESTIMATE_DELAY to estimate the delay of a packet in its buffer. ESTIMATE_DELAY works as follows (refer to box marked ALGORITHM ESTIMATE_DELAY): In Step 1, each node X maintains a separate queue of packets Q destined to a node Z sorted in decreasing order of creation times; this is the order in which the packets will be delivered when X meets Z in PROTOCOL RAPID. In Step 2 of ESTIMATE_DELAY, X computes the delivery delay distribution of packet i if delivered directly by X . In Step 3, X computes the minimum across all replicas of the corresponding delivery delay distributions; we note that the delivery time of i is the time until the first node delivers the packet. ESTIMATE_DELAY assumes that the meeting time distribution is the same as the inter-meeting time distribution.

The *Assumption 2* in ESTIMATE_DELAY is a simplifying independence assumption that does not hold in general. Consider Figure 2(a), an example showing the positions of packet replicas in the queues of different nodes. All packets have a common destination Z and each queue is sorted by $T(i)$. Assume that the transfer opportunities and packets are of unit-size.

In Figure 2(a), packet b may be delivered in two ways: (i) if W meets Z ; (ii) one of X and Y meets Z and then one of X and Y meet Z again. These delay dependencies can be represented using a dependency graph as illustrated in Fig 2(b); packets with the same letter and different indices are replicas. A vertex corresponds to a packet replica. An edge from one node to another indicates a dependency between the delays of the corresponding packets. Recall that M_{XY} is the random variable that represents the meeting time between X and Y .

ALGORITHM ESTIMATE_DELAY:

Node X storing a set of packets Q to destination Z performs the following steps to estimate the time until packet $i \in Q$ is delivered

- 1) X sorts all packets $i \in Q$ in the descending order of $T(i)$, time since i is created.

- a) Let $b(i)$ be the sum size of packets that precede packet i in the sorted list of X . Figure 1 illustrates a sorted buffer containing packet i .

- b) Let B be the expected transfer opportunity in bytes between X and Z . (For readability, we drop subscript X since we are only talking about one node; in general $b(i)$ and B are functions of the node). Node X locally computes B as a moving average of past transfers between X and Z .

- 2) *Assumption 1: Suppose only X delivers packets to Z with no further replication.*

Let $a_X(i)$ be the delay distribution of X delivering the packet. Under our assumption, X requires $\lceil b(i)/B \rceil$ meetings with Z to deliver i .

Let M be a distribution that models the inter-meeting times between nodes, and let $M_{X,Z}$ be the random variable that represents the time taken for X and Z to meet. We transform $M_{X,Z}$ to random variable $M'_{X,Z}$ that represents the time until X and Z meet $\lceil b(i)/B \rceil$ times. Then, by definition

$$a_X(i) = M'_{X,Z} \quad (4)$$

- 3) *Assumption 2: Suppose the k random variables $a_y(i), y \in [1, k]$ were independent, where k is the number of replicas of i .*

The probability of delivering i within time t is the minimum of the k random variables $a_y(i), y \in [1, k]$. This probability is:

$$P(a(i) < t) = 1 - \prod_{y=1}^k (1 - P(a_y(i) < t)) \quad (5)$$

- a) Accordingly: $A(i) = E[a(i)] \quad (6)$

ESTIMATE_DELAY ignores all the non-vertical dependencies. For example, it estimates b 's delivery time distribution as

$$\min(M_{WZ}, M_{XZ} + M_{XZ}, M_{YZ} + M_{YZ}),$$

whereas the distribution is actually

$$\min(M_{WZ}, \min(M_{XZ}, M_{YZ}) + \min(M_{XZ}, M_{YZ})).$$

Estimating delays without ignoring the non-vertical dependencies is challenging. Using a simplifying assumption that the transfer opportunities and packets are unit-sized, we design algorithm DAG_DELAY (described in a Technical report [citerapid-tr](#)), that estimates the expected delay by taking into account non-vertical dependencies. Although DAG_DELAY




Fig. 1. Position of packet i in a queue of packets destined to Z .




Fig. 2. Delay dependencies between packets destined to Z buffered in different nodes.

is of theoretical interest, it cannot be implemented in practice because DAG_DELAY assumes that — (i) the transfer opportunity size is exactly equal to the size of a packet. This assumption is fundamental for the design of DAG_DELAY and (ii) nodes have a global view of the system.

In general, ignoring non-vertical edges can arbitrarily inflate delay estimates for some pathological cases (detailed in a Technical report [3]). However, we find that ESTIMATE_DELAY works well in practice, and is simple and does not require a global view of the system.

1) *Estimating delays when transfer opportunities are exponentially distributed:* We walk through the distributed implementation of ESTIMATE_DELAY for a scenario where the inter-meeting time between nodes is exponentially distributed. Assume that the mean meeting time between nodes is $\frac{1}{\lambda}$. In the absence of bandwidth restrictions, the expected delivery delay when there are k replicas is the mean meeting time divided by k , i.e., $P(a(i) < t) = 1 - e^{-k\lambda t}$ and $A(i) = \frac{1}{k\lambda}$. (Note that the minimum of k i.i.d. exponentials is also an exponential with mean $\frac{1}{k}$ of the mean of the i.i.d. exponentials [9].)

When transfer opportunities are limited, the expected delay depends on the packet's position in the nodes' buffers. In Step 2 of ESTIMATE_DELAY, the node estimates the number of times it needs to meet the destination to deliver a packet as a function of $\lceil b(i)/B \rceil$. According to our exponential meeting time assumption, the time for some node X to meet the destination $\lceil b(i)/B \rceil$ times is described by a gamma distribution with mean $\frac{1}{\lambda} \cdot \lceil b(i)/B \rceil$.

If packet i is replicated at k nodes, Step 3 computes the delay distribution $a(i)$ as the minimum of k gamma variables. We do not know of a closed form expression for the minimum of gamma variables. Instead, if we assume that the time taken for a node to meet the destination $b(i)/B$ times is exponential with the same mean $\frac{1}{\lambda} \cdot \lceil b(i)/B \rceil$. We can then estimate $a(i)$ as the minimum of k exponentials.

Let $n_1(i), n_2(i), \dots, n_k(i)$ be the number of times each of

the k nodes respectively needs to meet the destination to deliver i directly. Then $A(i)$ is computed as:

$$P(a(i) < t) = 1 - e^{-(\frac{\lambda}{n_1(i)} + \frac{\lambda}{n_2(i)} + \dots + \frac{\lambda}{n_k(i)})t} \quad (7)$$

$$A(i) = \frac{1}{\frac{\lambda}{n_1(i)} + \frac{\lambda}{n_2(i)} + \dots + \frac{\lambda}{n_k(i)}} \quad (8)$$

When the meeting time distributions between nodes are non-uniform, say with means $\frac{1}{\lambda_1}, \frac{1}{\lambda_2} \dots \frac{1}{\lambda_k}$ respectively, then $A(i) = (\frac{\lambda_1}{n_1(i)} + \frac{\lambda_2}{n_2(i)} + \dots + \frac{\lambda_k}{n_k(i)})^{-1}$.

2) *Estimating delays when transfer opportunity distribution is unknown*: To implement RAPID on the DieselNet testbed, we adapt Eq. 8 to scenarios where the transfer opportunities are not exponentially distributed. First, to estimate mean inter-node meeting times in the DieselNet testbed, every node tabulates the average time to meet every other node based on past meeting times. Nodes exchange this table as part of metadata exchanges (Step 1 in PROTOCOL RAPID). A node combines the metadata into a meeting-time adjacency matrix and the information is updated after each transfer opportunity. The matrix contains the expected time for two nodes to meet directly, calculated as the average of past meetings.

Node X estimates $E(M_{XZ})$, the expected time to meet Z , using the meeting-time matrix. $E(M_{XZ})$ is estimated as the expected time taken for X to meet Z in at most h hops. (Unlike uniform exponential mobility models, some nodes in the trace never meet directly.) For example, if X meets Z via an intermediary Y , the expected meeting time is the expected time for X to meet Y and then Y to meet Z in 2 hops. In our implementation we restrict $h = 3$. When two nodes never meet, even via three intermediate nodes, we set the expected inter-meeting time to infinity. Several DTN routing protocols [7], [21], [8] use similar techniques to estimate meeting probability among peers.

RAPID estimates expected meeting times by taking into account transitive meetings. However, our delivery estimation (described in ESTIMATE_DELAY) assumes that nodes do not make additional replicas. This disconnect is because, in DieselNet, only few buses meet directly, and the pair-wise meeting times between several bus pairs is infinity. We take into account transitive meetings when two buses do not meet directly, to increase the number of potential forwarders.

Let replicas of packet i destined to Z reside at nodes X_1, \dots, X_k . Since we do not know the meeting time distributions, we simply assume they are exponentially distributed. Then from Eq. 8, the expected delay to deliver i is

$$A(i) = \left[\sum_{j=1}^k \frac{1}{E(M_{X_j Z}) \cdot n_j(i)} \right]^{-1} \quad (9)$$

We use an exponential distribution because bus meeting times in the testbed are difficult to model. Buses change routes several times in one day, the inter-bus meeting distribution is noisy, and we found them hard to model even using mixture models. Approximating meeting times as exponentially distributed makes delay estimates easy to compute and performs well in practice.

B. Control channel

Previous studies [18] have shown that as nodes have the benefit of more information about global system state using oracles, they can make significantly better routing decisions. We extend this idea to practical DTNs where no oracle is available. RAPID nodes gather knowledge about the global system state by disseminating metadata using a fraction of the transfer opportunity.

RAPID uses an in-band *control channel* to exchange acknowledgments for delivered packets as well as metadata about every packet learnt from past exchanges. For each encountered packet i , RAPID maintains a list of nodes that carry the replica of i , and for each replica, an estimated time for direct delivery. Metadata for delivered packets is deleted when an ack is received.

For efficiency, a RAPID node maintains the time of last metadata exchange with its peers. The node only sends information about packets whose information changed since the last exchange, which reduces the size of the exchange considerably. A RAPID node sends the following information on encountering a peer: (i) Average size of past transfer opportunities; (ii) Expected meeting times with nodes; (iii) Acks; (iv) For each of its own packets, the updated delivery delay estimate based on current buffer state; (v) Delivery delay of other packets if modified since last exchange.

When using the control channel, nodes have only an imperfect view of the system. The propagated information may be stale due to changes in number of replicas, changes in delivery delays, or if the packet is delivered but acknowledgments have not propagated. Nevertheless, our experiments confirm that (i) this inaccurate information is sufficient for RAPID to achieve significant performance gains over existing protocols and (ii) the overhead of metadata itself is not significant.

V. THE CASE FOR A HEURISTIC APPROACH

Any DTN routing algorithm has to deal with two uncertainties regarding the future: unpredictable meeting schedule and unpredictable workload. RAPID is a local algorithm that routes packets based on the marginal utility *heuristic* in the face of these uncertainties. In this section, we show two fundamental reasons that make the case for a heuristic approach to DTN routing. First, we prove that computing optimal solutions is hard even with complete knowledge about the environment. Second, we prove that the presence of even one of the two uncertainties rule out provably efficient online routing algorithms.

A. Computational Hardness of the DTN Routing Problem

THEOREM 2: Given complete knowledge of node meetings and the packet workload *a priori*, computing a routing schedule that is optimal with respect to the number of packets delivered is NP-hard and has a lower bound of $\Omega(n^{1/2-\epsilon})$ on the approximation ratio.

Proof: Consider a DTN routing problem with n nodes that have complete knowledge of node meetings and workload *a priori*. The input to the DTN problem is the set of nodes $1, \dots, n$; a series of transfer opportunities $\{(u_1, v_1, s_1, t_1), (u_2, v_2, s_2, t_2), \dots\}$ such that $u_i, v_i \in [1, n]$, s_i is the size of the transfer opportunity, and t_i is the time

of meeting; and a packet workload $\{p_1, p_2, \dots, p_s\}$, where $p_i = (u'_i, v'_i, s'_i, t'_i)$, where $u', v' \in [1, n]$ are the source and destination, s' the size, and t' the time of creation of the packet, respectively. The goal of a DTN routing algorithm is to compute a feasible schedule of packet transfers, where *feasible* means that the total size of transferred packets in any transfer opportunity is less than the size of the transfer opportunity.

The decision version $O_{n,k}$ of this problem is: Given a DTN with n nodes such that nodes have complete knowledge of transfer opportunities and the packet workload, is there a feasible schedule that delivers at least k packets?

LEMMA 1: $O(n, k)$ is NP-hard.

Proof: We show that $O(n, k)$ is a NP-hard problem using a polynomial-time reduction from the edge-disjoint path (EDP) problem for a directed acyclic graph (DAG) to $O(n, k)$. The EDP problem for a DAG is known to be NP-hard [11].

The decision version of EDP problem is: Given a DAG $G = (V, E)$, where $|V| = n$, $E \in V \times V$: $e_i = (u_i, v_i) \in E$, if e_i is incident on u_i and v_i and direction is from u_i to v_i . If given source-destination pairs $\{(s_1, t_1), (s_2, t_2) \dots (s_s, t_s)\}$, do a set of edge-disjoint paths $\{c_1, c_2 \dots c_k\}$ exist, such that c_i is a path between s_i and t_i , where $1 \leq i \leq k$.

Given an instance of the EDP problem, we generate a DTN problem $O(n, k)$ as follows:

As the first step, we topologically order the edges in G , which is possible given G is a DAG. The topological sorting can be performed in polynomial-time.

Next, we label edges using natural numbers with any function $l : E \rightarrow \mathbb{N}$ such that if $e_i = (u_i, u_j)$ and $e_j = (u_j, u_k)$, then $l(e_i) < l(e_j)$. There are many ways to define such a function l . One algorithm is:

- 1) label = 0
- 2) For each vertex v in the decreasing order of the topological sort,
 - a) Choose unlabeled edge $e = (v, x) : x \in V$,
 - b) label = label + 1
 - c) Label e ; $l(e) = \text{label}$.

Since vertices are topologically sorted, if $e_i = (u_i, u_j)$ then $u_i < u_j$. Since the algorithm labels all edges with source u_i before it labels edges with source u_j , if $e_j = (u_j, u_k)$, then $l(e_i) < l(e_j)$.

Given a G , we define a DTN routing problem by mapping V to the nodes $(1, \dots, n)$ in the DTN. The edge $(e = \{u, v\} : u, v \in V)$ is mapped to the transfer opportunity $(u, v, 1, l(e))$, assuming transfer opportunities are unit-sized. Source and destination pairs $\{(s_1, t_1), (s_2, t_2), \dots, (s_m, t_m)\}$ are mapped to packets $\{p_1, p_2, \dots, p_m\}$, where $p_i = (s_i, t_i, 1, 0)$. In other words, packet p is created between the corresponding source-destination pair at time 0 and with unit size. A path in graph G is a valid route in the DTN because the edges on a path are transformed to transfer opportunities of increasing time steps. Moreover, a transfer opportunity can be used to send no more than one packet because all opportunities are unit-sized. If we solve the DTN routing problem of delivering k packets, then there exists k edge-disjoint paths in graph G , or in other words we can solve the EDP problem. Similarly, if the EDP problem has a solution consisting of k edge-disjoint

paths in G , at least k packets can be delivered using the set of transfer opportunities represented by each path. Using the above polynomial-time reduction, we show that a solution to EDP exists if and only if a solution to $O(n, k)$ exists. Thus, $O(n, k)$ is NP-hard. ■

COROLLARY 1: The DTN routing problem has a lower bound of $\Omega(n^{1/2-\epsilon})$ on the approximation ratio.

Proof: The reduction given above is a true reduction in the following sense: each successfully delivered DTN packet corresponds to an edge-disjoint path and vice-versa. Thus, the optimal solution for one exactly corresponds to an optimal solution for the other. Therefore, this reduction is an L-reduction [25]. Consequently, the lower bound $\Omega(n^{1/2-\epsilon})$ known for the hardness of approximating the EDP problem [15] holds for the DTN routing problem as well. ■

Hence, Theorem 2. ■

The hardness results naturally extend to the average delay metric for both the online as well as computationally limited algorithms.

B. Competitive Hardness of Online DTN Routing




Fig. 3. DTN node meetings for Theorem V-B. Solid arrows represent node meetings known a priori to the online algorithm while dotted arrows represent meetings revealed subsequently by an offline adversary.

Let ALG be any deterministic online DTN routing algorithm with unlimited computational power.

THEOREM 1(a). *If ALG has complete knowledge of the workload, but not of the schedule of node meetings, then ALG is $\Omega(n)$ -competitive with an offline adversary with respect to the fraction of packets delivered, where n is the number of packets in the workload.*

Proof: We prove the theorem by constructing an offline adversary, ADV, that incrementally generates a node meeting schedule after observing the actions of ALG at each step. We show how ADV can construct a node meeting schedule such that ADV can deliver all packets while ALG, without prior knowledge of node meetings, can deliver at most 1 packet.

Consider a DTN as illustrated in Fig. 3, where $P = \{p_1, p_2, \dots, p_n\}$ denotes a set of unit-sized packets; $U = \{u_1, u_2, \dots, u_n\}$ denotes a set of intermediate nodes; and $V = \{v_1, v_2, \dots, v_n\}$ denotes a set of nodes to which the packets are respectively destined, i.e. p_i is destined to v_i

for all $i \in [1, n]$. The following procedure describes ADV's actions given ALG as input.

PROCEDURE FOR ADV:

- **Step 1:** ADV generates a set of node meetings involving unit-size transfer opportunities at time $t = 0$ between A and each of the intermediate nodes u_1, \dots, u_n respectively (refer to Figure 3).
- **Step 2:** At time $t_1 > 0$, ADV observes the set of transfers X made by ALG. Without loss of generality, $X : P \rightarrow U$ is represented as a (one-to-many) mapping where $X(p_i)$ is the set of intermediate nodes $(u_1, u_2 \dots u_n)$ to which ALG replicates packet p_i .
- **Step 3:** ADV generates the next set of node meetings $(u_1, Y(u_1)), (u_2, Y(u_2)), \dots, (u_n, Y(u_n))$ at time t_1 , where $Y : U \rightarrow V$ is a bijective mapping from the set of intermediate nodes to the destination nodes $v_1, v_2, \dots v_n$.

ADV uses the following procedure to generate the mapping Y given X in Step 3.

PROCEDURE GENERATE_Y(X):

- 1) Initialize $Y(p_i)$ to null for all $i \in [1, n]$;
- 2) for each $i \in [1, n]$ do
- 3) if $\exists j : u_j \notin X(p_i)$ and $Y(u_j) = \text{null}$, then
- 4) Map $Y(u_j) \rightarrow v_i$ for the smallest such j ;
- 5) else
- 6) Pick a $j : Y(u_j) = \text{null}$, and map $Y(u_j) \rightarrow v_i$
- 7) endif

LEMMA 2: ADV executes Line 6 in GENERATE_Y(X) at most once.

Proof:

We first note that the procedure is well defined at Line 6: each iteration of the main loop map exactly one node in U to a node in V , therefore a suitable j such that $Y(u_j) = \text{null}$ exists. Suppose ADV first executes Line 6 in the m 'th iteration. By inspection of the code, the condition in Line 3 is false, therefore each intermediate node $u_k, k \in [1, n]$, either belongs to $X(p_i)$ or is mapped to some destination node $Y(u_k) \neq \text{null}$. Since each of the $m - 1$ previous iterations must have executed Line 4 by assumption, exactly $m - 1$ nodes in U have been mapped to nodes in V . Therefore, each of the remaining $n - m + 1$ unmapped nodes must belong to $X(p_i)$ in order to falsify Line 3. Line 6 maps one of these to v_i leaving $n - m$ unmapped nodes. None of these $n - m$ nodes is contained in $X(p_k)$ for $k \in [m + 1, \dots, n]$. Thus, in each of the subsequent $n - m$ iterations, the condition in Line 3 evaluates to true. ■

LEMMA 3: The schedule of node meetings created by Y allows ALG to deliver at most one packet to its destination.

Proof:

For ALG to deliver any packet p_i successfully to its destination v_i , it must be the case that some node in $X(p_i)$ maps to v_i . Such a mapping could not have occurred in Line 3 by inspection of the code, so it must have occurred in Line 6. By Lemma 2, Line 6 is executed exactly once, so ALG can deliver at most one packet. ■

LEMMA 4: The schedule of node meetings created by Y allows ADV to deliver all packets to their respective destinations.

Proof: We first note that, by inspection of the code, Y is a bijective mapping: Line 4 and 6 map an unmapped node in U to v_i in iteration m and there are n such iterations. So, ADV can route p_i by sending it $Y^{-1}(v_i)$ and subsequently to v_i . ■

Theorem 1(a) follows directly from Lemmas 3 and 4. ■

COROLLARY 2: ALG can be arbitrarily far from ADV with respect to average delivery delay.

Proof: The average delivery delay is unbounded for ALG because of undelivered packets in the construction above while it is finite for ADV. If we assume that that ALG can eventually deliver all packets after a long time T (say, because all nodes connect to a well-connected wired network at the end of the day), then ALG is $\Omega(T)$ -competitive with respect to average delivery delay using the same construction as above. ■

We remark that it is unnecessary in the construction above for the two sets of n node meetings to occur simultaneously at $t = 0$ and $t = t_1$, respectively. The construction can be easily modified to not involve any concurrent node meetings.

THEOREM 1(b). *If ALG has complete knowledge of the meeting schedule, but not of the packet workload, then ALG can deliver at most a third of the packets delivered by an optimal offline adversary.*

Proof: We prove the theorem by constructing a procedure for ADV to incrementally generate a packet workload by observing ALG's transfers at each step. As before, we only need unit-sized transfer opportunities and packets for the construction.

Consider the basic DTN "gadget" shown in Fig. 4(a) involving just six node meetings. The node meetings are known in advance and occur at times T_1 and $T_2 > T_1$ respectively. The workload consists of just two packets $P = \{p_1, p_2\}$ destined to v_1 and v_2 , respectively.

LEMMA 5: ADV can use the basic gadget to force ALG to drop half the packets while itself delivering all packets.

Proof: The procedure for ADV is as follows. If ALG transfers p_1 to v'_1 and p_2 to v'_2 , then ADV generates two more packets: p'_2 at v'_1 destined to v_2 and p'_1 at v'_2 destined to v_1 . ALG is forced to drop one of the two packets at both v'_1 and v'_2 . ADV can deliver all four packets by transferring p_1 and p_2 to v'_2 and v'_1 respectively at time T_1 , which is the exact opposite of ALG's choice.

If ALG instead chooses to transfer p_1 to v'_2 and p_2 to v'_1 , ADV chooses the opposite strategy.

If ALG chooses to replicate one of the two packets in both transfer opportunities at time T_1 while dropping the other packet, ADV simply deliver both packets. Hence the lemma. ■

Next, we extend the basic gadget to show that ALG can deliver at most a third of the packets while ADV delivers all packets. The corresponding construction is shown in Figure 4(b).

The construction used by ADV composes the basic gadget repeatedly for a depth of 2. In this construction, ADV can force ALG to drop 2/5th of the packet while ADV delivers all packets. We provide the formal argument in a technical

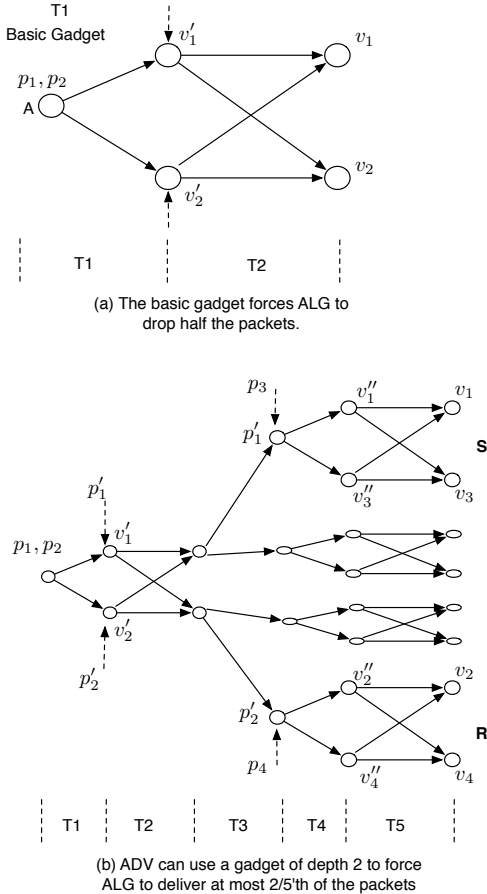


Fig. 4. DTN construction for Theorem V-B. Solid arrows represent node meetings known a priori to ALG while vertical dotted arrows represent packets created by ADV at the corresponding node.

report [3] in the interest of space. Similarly, by creating a gadget of depth 3, we can show that ADV can force ALG to deliver at most $4/11$ 'th of the packets. Effectively, each new basic gadget introduces 3 more packets and forces ALG to drop 2 more packets. In particular, with a gadget of depth i , ADV can limit ALG's delivery rate to $i/(3i - 1)$. Thus, by composing a sufficiently large number of basic gadgets, ADV can limit the delivery rate of ALG to a value close to $1/3$.

Hence, Theorem 1(b). ■

VI. IMPLEMENTATION ON A VEHICULAR DTN TESTBED

We implemented and deployed RAPID on our vehicular DTN testbed, DieselNet [7] (<http://prisms.cs.umass.edu/dome>), consisting of 40 buses, of which a subset is on the road each day. The routing protocol implementation is a first step towards deploying realistic DTN applications on the testbed. In addition, the deployment allows us to study the effect of certain events that are not perfectly modeled in the simulation of our routing protocol. These events include delays caused by computation, wireless channel interference, and operating system delays.

Each bus in DieselNet carries a small-form desktop computer, 40 GB of storage, and a GPS device. The buses operate a 802.11b radio that scans for other buses 10 times a second and an 802.11b access point (AP) that accepts incoming connections.

Once a bus is found, a connection is created to the remote AP. (It is likely that the remote bus then creates a connection to the discovered AP, which our software merges into one connection event.) The connection lasts until the radios are out of range. Burgess et al. [7] describes the DieselNet testbed in more detail.

A. Deployment

Buses in DieselNet send messages using PROTOCOL RAPID in Section III, computing the metadata as described in Section IV-B. We generated packets of size 1 KB periodically on each bus with an exponential inter-arrival time. The destinations of the packets included only buses that were scheduled to be on the road, which avoided creation of many packets that could never be delivered. We did not provide the buses information about the location or route of other buses on the road. We set the default packet generation rate to 4 packets per hour generated by each bus for every other bus on the road; since the number of buses on the road at any time varies, this is the simplest way to express load. For example, when 20 buses are on the road, the default rate is 1,520 packets per hour.

During the experiments, the buses logged packet generation, packet delivery, delivery delay, meta-data size, and the total size of the transfer opportunity. Buses transferred random data after all routing was complete in order to measure the capacity and duration of each transfer opportunity. The logs were periodically uploaded to a central server using open Internet APs found on the road.

B. Performance of deployed RAPID

We measured the routing performance of RAPID on the buses from Feb 6, 2007 until May 14, 2007¹. The measurements are tabulated in Table III. We exclude holidays and weekends since almost no buses were on the road, leaving 58 days of experiments. RAPID delivered 88% of packets with an average delivery delay of about 91 minutes. We also note that overhead due to meta-data accounts for less than 0.2% of the total available bandwidth and less than 1.7% of the data transmitted.

C. Validating trace-driven simulator

In the next section, we evaluate RAPID using a trace-driven simulator. The simulator takes as input a schedule of node meetings, the bandwidth available at each meeting, and a routing algorithm. We validated our simulator by comparing simulation results against the 58-days of measurements from the deployment. In the simulator, we generate packets under the same assumptions as the deployment, using the same parameters for exponentially distributed inter-arrival times.

Figure 5 shows the average delay characteristics of the real system and the simulator. Delays measured using the simulator were averaged over the 30 runs and the error-bars show a 95% confidence interval. From those results and further analysis, we find with 95% confidence that the simulator results are within 1% of the implementation measurement of average delay. The close correlation between system measurement and simulation increases our confidence in the accuracy of the simulator.

¹The traces are available at <http://traces.cs.umass.edu>.

Avg. buses scheduled per day	19
Avg. total bytes transferred per day	261.4 MB
Avg. number of meetings per day	147.5
Percentage delivered per day	88%
Avg. packet delivery delay	91.7 min
Meta-data size/ bandwidth	0.002
Meta-data size/ data size	0.017

TABLE III
DEPLOYMENT OF RAPID: AVERAGE DAILY STATISTICS

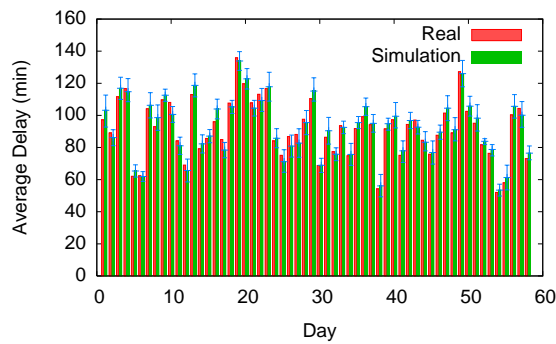


Fig. 5. **Trace:** Average delay for 58 days of RAPID real deployment compared to simulation of RAPID using traces

VII. EVALUATION

The goal of our evaluation is to show that, unlike existing work, RAPID can improve performance for customizable metrics. We evaluate RAPID using three metrics: minimize maximum delay, minimize average delay, and minimize missed deadlines. In all cases, we found that RAPID significantly outperforms existing protocols and also performs close to optimal for small workloads.

A. Experimental setup

Our evaluations are based on a custom event-driven simulator, as described in the previous section. The meeting times between buses in these experiments are not known *a priori*. All values used by RAPID, including average meeting times, are learned during the experiment.

We compare RAPID to five other routing protocols: *MaxProp* [7], *Spray and Wait* [29], *Prophet* [21], *Random*, and *Optimal*. In all experiments, we include the cost of RAPID’s in-band control channel for exchanging metadata.

MaxProp operates in a storage- and bandwidth-constrained environment, allows packet replication, and leverages delivery notifications to purge old replicas; of recent related work, it is closest to RAPID’s objectives. *Random* replicates randomly chosen packets for the duration of the transfer opportunity. *Spray and Wait* restricts the number of replications of a packets to L , where L is calculated based on the number of nodes in the network. For our simulations, we implemented the binary *Spray and Wait* and set² $L = 12$. We implemented *Prophet* with

²We set this value based on consultation with authors and using LEMMA 4.3 in [29] with $a = 4$.

	Exponential/ Power law	Trace-driven
Number of nodes	20	max of 40
Buffer size	100 KB	40 GB
Transfer opp. size	100 KB	given by trace
Duration	15 min	19 hours each trace
Size of a packet	10 KB	10 KB
Packet generation rate	50 sec mean	1 hour
Delivery deadline	20 sec	2.7 hours

TABLE IV
EXPERIMENT PARAMETERS

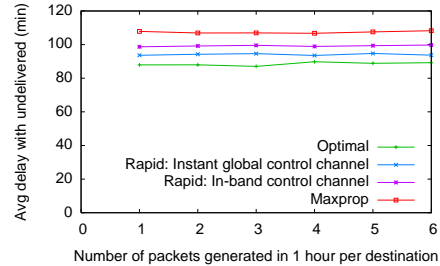


Fig. 15. (Trace) Comparison with *Optimal*: Average delay of RAPID is within 10% of *Optimal* for small loads

parameters $P_{init} = 0.75$, $\beta = 0.25$ and $\gamma = 0.98$ (parameters based on values used in [21]).

We also perform experiments where mobility is modeled using a synthetic distribution – in this work we consider exponential and power law distribution. Previous studies [10], [20] have suggested that DTNs among people have a skewed, power law inter-meeting time distribution. The default parameters used for all the experiments are tabulated in Table IV. The parameters for the synthetic mobility model is different from the trace-driven model because the performance between the two models are not comparable.

Each data point is averaged over 10 runs; in the case of trace-driven results, the results are averaged over 58 traces. Each of the 58 days is a separate experiment. In other words, packets that are not delivered by the end of the day are lost. In all experiments, *MaxProp*, RAPID and *Spray and Wait* performed significantly better than *Prophet*, and the latter is not shown in the graphs for clarity.

B. Results based on testbed traces

1) *Comparison with existing routing protocols:* Our experiments show that RAPID consistently outperforms *MaxProp*, *Spray and Wait* and *Random*. We increased the load in the system up to 40 packets per hour per destination, when *Random* delivers less than 50% of the packets.

Figure 6 shows the average delay of delivered packets using the four protocols for varying loads when RAPID’s routing metric is set to minimize average delay (Eq. 1). When using RAPID, the average delay of delivered packets are significantly lower than *MaxProp*, *Spray and Wait* and *Random*. Moreover, RAPID also consistently delivers a greater fraction of packets as shown in Figure 7.




Fig. 6. (Trace) Average Delay: RAPID has up to 20% lower delay than *MaxProp* and up to 35% lower delay than *Random*




Fig. 7. (Trace) Delivery Rate: RAPID delivers up to 14% more than *MaxProp*, 28% than *Spray and Wait* and 45% than *Random*




Fig. 8. (Trace) Max Delay: Maximum delay of RAPID is up to 90 min lower than *MaxProp*, *Spray and Wait*, and *Random*




Fig. 9. (Trace) Delivery within deadline: RAPID delivers up to 21% more than *MaxProp*, 24% than *Spray and Wait*, 28% than *Random*




Fig. 10. (Trace) Control channel benefit: Average delay performance improves as more metadata is allowed to be exchanged




Fig. 11. (Trace) Channel utilization: As load increases, delivery rate decreases to 65% but channel utilization is only about 35%

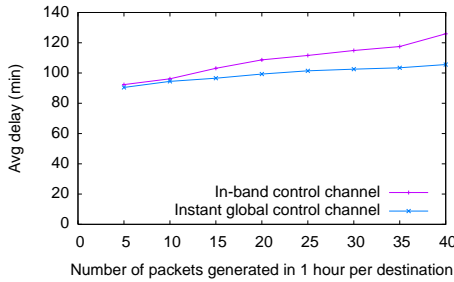


Fig. 12. (Trace) Global channel: Average delay of RAPID decreases by up to 20 minutes using instant global control channel

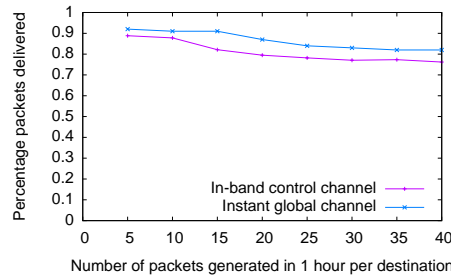


Fig. 13. (Trace) Global channel: Delivery rate increases by up to 12% using an instant global control channel, for the average delay metric




Fig. 14. (Trace) Global channel: Packets delivered within deadline increases by about 15% using instant global control channel

Figure 8 shows RAPID's performance when the routing metric is set to minimize maximum delay (Eq. 3) and similarly Figure 9 shows results when the metric is set to maximize the number of packets delivered within a deadline (Eq. 2).

We note that among *MaxProp*, *Spray and Wait* and *Random*, *MaxProp* delivers the most number of packets, but *Spray and Wait* has marginally lower average delay than *MaxProp*. RAPID significantly outperforms the three protocol for all metrics because of its intentional design.

Standard deviation and similar measures of variance are not appropriate for comparing the mean delays as each bus takes a different geographic route. So, we performed a paired *t*-test [9] to compare the average delay of every source-destination pair using RAPID to the average delay of the same source-destination pair using *MaxProp* (the second best performing protocol). In our tests, we found *p*-values always less than 0.0005, indicating the differences between the means reported in these figures are statistically significant.

In a separate experiment (not shown in figure), we find that the number of replications per delivery made by RAPID is 5.2, for a load of 5 packets per hour per destination. For the same load, the number of replications per delivery made by *Random* is 3.5 and *Spray and Wait* is 4.2. We note that we only consider the number of replications for packets that are delivered, and RAPID is set to optimize the average delay metric. Even though it seems that RAPID replicates more aggressively to deliver more packets, RAPID only replicates when bandwidth is available. For example, when the load is increased to 15 packets per hour per destination, the number of replications per delivery made by RAPID reduced to 4.3.

2) *Metadata exchange*: We allow RAPID to use as much bandwidth at the start of a transfer opportunity for exchanging metadata as it requires. To see if this approach was wasteful or beneficial, we performed experiments where we limited the total metadata exchanged. Figure 10 shows the average delay performance of RAPID when metadata is limited as a

percentage of the total bandwidth. The average delay metric shown here includes the delay for undelivered packets. When a packet is undelivered, it is assumed to be delivered at the end of the day.

The results show that performance increases as the limit is removed and that the best performance results when there is no restriction on metadata at all. The performance of RAPID with complete metadata exchange improves by 20% compared to when no metadata is exchanged. The metadata in this experiment is represented as a percentage of available bandwidth.

In the next experiment, we analyze total metadata as a percentage of data. In particular, we increase the load to 75 packets per destination per hour to analyze the trend in terms of bandwidth utilization, delivery rate and metadata. Figure 11 shows this trend as load increases. The bandwidth utilization is about 35% for the load of 75 packets per hour per destination, while delivery rate is only about 65%. This suggests that the performance drops even though the network is under-utilized, and it is because of the bottleneck links in the network. The available bandwidth varies significantly across transfer opportunities in our bus traces [7].

We also observe that metadata increases to about 4% of data for high loads. This is an order of magnitude higher than the metadata observed as a fraction of bandwidth, again because of the poor channel utilization. The average metadata exchange per contact is proportional to the load and the channel utilization.

RAPID uses more information to improve routing performance. Although the result is intuitive, RAPID uses the additional information to compute packet utilities accurately and in-turn replicate packets intentionally. In contrast, *Spray and Wait* or *Random* cannot use additional information even if available, and *MaxProp* uses additional information only to remove delivered packets [7]. Further, collecting the additional information does not incur a huge overhead in RAPID. The metadata overhead reduces even further with increasing packet size. For example, moving from 1-KB to 10-KB packets reduces RAPID's metadata overhead by an order of magnitude.

There are several scenarios where metadata exchange needs to be limited. For example, when transfer opportunities sizes are much smaller than the number of packets, exchanging all metadata during a transfer opportunity may affect performance. Similarly, since RAPID is a link-state routing protocol, it scales only as well as a link-state protocol. As the network size increases, a node may need to limit the state information it maintains as well as the amount of metadata exchanged. The issue of limiting metadata exchange according to the network scenario will be addressed as part of future work.

3) *Hybrid DTN with thin continuous connectivity*: In this section, we compare the performance of RAPID using an instant *global* control channel for exchanging metadata as opposed to the default (delayed) *in-band* control channel.

Figure 12 shows the average delay of RAPID when using an in-band control channel compared to a global channel. We observe that the average delay of delivered packets decreases by up to 20 minutes when using a global channel. For the same experiments, the delivery rate when using an instant global

channel increases by up to 12% (shown in Figure 13). Similarly, Figure 14 shows that the percentage packets delivered within a deadline increases by an average of 20% using a global channel. This observation suggests that RAPID's performance can benefit further by using more control information.

One interpretation of the global channel is the use of RAPID as a hybrid DTN where all control traffic goes over a low-bandwidth, long-range radio such as XTend [4]. Since XTend radios support a data rate of about 1 Kbps for a range of 1 mile, the radios cannot be used to deliver data packets when the incoming rate is high or packet sizes are large. A hybrid DTN will use a high-cost, low-bandwidth channel for control whenever available and low-cost high-bandwidth delayed channel for data. In our experiments, we assumed that the global channel is instant. While this may not be feasible in practice, the results give an upper bound on RAPID's performance when accurate channel information is available.

C. Results compared with Optimal

We compare RAPID to *Optimal*, which is an upper bound on the performance. To obtain the optimal delay, we formulate the DTN routing problem as an Integer Linear Program (ILP) optimization problem when the meeting times between nodes are precisely known (details in a Technical report [3]) and solve the problem using a CPLEX solver [12]. Because the problem grows in complexity with the number of packets, these results are limited to only 6 packets per hour per destination. The ILP objective function minimizes delay of all packets, where the delay of undelivered packets is set to time the packet spent in the system. Accordingly, we add the delay of undelivered packets when presenting the results for RAPID and *MaxProp*.

Figure 15 presents the average delay performance of *Optimal*, RAPID, and *MaxProp*. We observe that for small loads, the performance of RAPID using the in-band control channel is within 10% of the optimum performance, while using *MaxProp* the delays are about 22% from the optimal. RAPID using a global channel performs within 6% of optimal.

D. Results from synthetic mobility models

Next, we use an exponential and power law mobility model to compare the performance of RAPID to *MaxProp*, *Random*, and *Spray and Wait*. When mobility is modeled using power law, two nodes meet with an exponential inter-meeting time, but the mean of the exponential distribution is determined by the popularity of the nodes. For the 20 nodes, we randomly set a popularity value of 1 to 20, with 1 being most popular.

We set the mean meeting time for both mobility distribution to 30 seconds. For the power law mobility model, the meeting time is skewed from 30 seconds according to the node's popularity. All other parameters for exponential and power law are identical.

1) *Powerlaw mobility model, increasing load*: Figure 16 shows the average delay for packets to be delivered (i.e., RAPID is set to use Eq. 1 as a metric). The average delay of packets quickly increase to 20 seconds as load increases in the case of *MaxProp*, *Spray and Wait* and *Random*. In comparison,

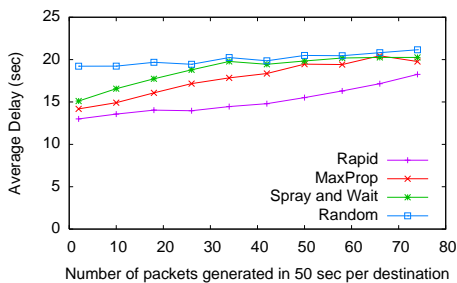


Fig. 16. (Powerlaw) Avg Delay: RAPID reduces delay by about 20% compared to *MaxProp*, and 23% than *Spray and Wait* and 25% than *Random*

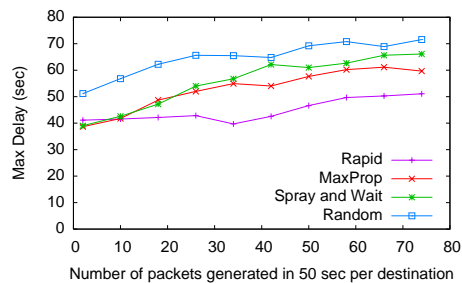


Fig. 17. (Powerlaw) Max delay: RAPID's max delay is about 30% lower than *MaxProp*, 35% lower than *Spray and Wait* and 45% lower than *Random*




Fig. 18. (Powerlaw) Delivery Deadline: RAPID delivers about 20% more packets within deadline when buffer size is constrained, compared to *MaxProp*, and 45% more packets compared to *Spray and Wait* and *Random*

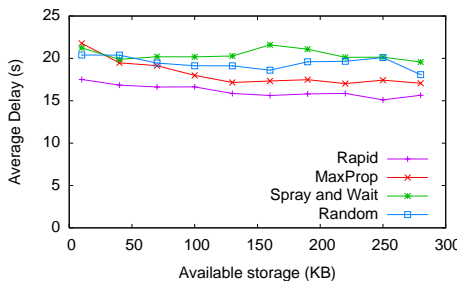


Fig. 19. (Powerlaw) Avg Delay with constrained buffer: RAPID reduces average delay by about 23% when buffer size is constrained compared to *MaxProp*, *Spray and Wait* and *Random*




Fig. 20. (Powerlaw) Max delay with constrained buffer: RAPID's max delay is about 22% lower than *MaxProp*, 35% lower than *Spray and Wait* and 38% lower than *Random* when buffer is constrained

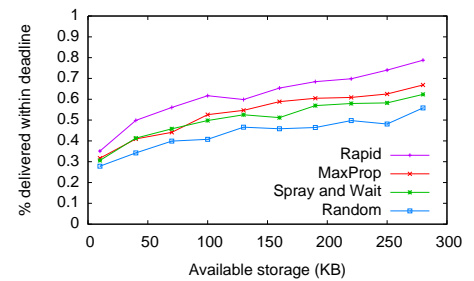


Fig. 21. (Powerlaw) Delivery Deadline with constrained buffer: RAPID delivers about 20% more packets within deadline when buffer size is constrained compared to *MaxProp*, and 45% more than *Spray and Wait* and *Random*

RAPID's delay does not increase rapidly with increasing load, and is on an average 20% lower than all the three protocols.

Figure 17 shows the performance with respect to minimizing the maximum delay of packets (using Eq. 3 as a metric). RAPID reduces maximum delay by an average of 30% compared to the other protocols. For both the traces and the synthetic mobility, the performance of RAPID is significantly higher than *MaxProp*, *Spray and Wait*, and *Random* for the maximum delay metric. The reason is *MaxProp* prioritizes new packets; older, undelivered packets will not see service as load increases. Similarly, *Spray and Wait* does not give preference to older packets. However, RAPID specifically prioritizes older packets to reduce maximum delay.

We observe similar trends in Figure. 18, that shows the performance of the different routing protocols with respect to maximizing the number of packet delivered within an average deadline of 20 sec (RAPID uses Eq. 2).

2) Powerlaw mobility model: decreasing storage constraint:

In this set of experiments, we varied available storage from 10 KB to 280 KB and compared the performance of the four routing protocols. We fixed the load to 20 packets per destination and generated packets with a inter-arrival time of 50 seconds.

Figure 19 shows how the average delay of all four protocols vary with increase storage availability. RAPID is able to maintain low delays even when only 10 KB space is available at each node. In comparison, *MaxProp*, *Spray and Wait* and *Random* have an average 23% higher delay.

Figure 20 shows a similar performance trend in terms of

minimizing maximum delay. Similar to other experiments, the difference in performance between RAPID and the other three protocols is more marked for the maximum delay metric.

Figure 21 shows how constrained buffers affect the delivery deadline metric. When storage is restricted, *MaxProp* deletes packets that are replicated most number of times, while *Spray and Wait* and *Random* deletes packets randomly. RAPID, when set to maximizing number of packets delivered within a deadline, deletes packets that are most likely to miss the deadline. RAPID is able to best manage limited buffers to deliver packets within a deadline and improves delivery performance by 12% compared to the second-best performing protocol. These experiments suggest that RAPID's utility-driven approach adapts well to storage restrictions as well. We observed similar trends for increasing storage restrictions when using exponential mobility model (not shown in figure).

VIII. CONCLUSIONS

Previous work in DTN routing protocols has seen only incidental performance improvement from various routing mechanisms and protocol design choices. In contrast, we have proposed a routing protocol for DTNs that intentionally optimizes a specific routing metric by treating DTN routing as a resource allocation problem. Although our approach is heuristic-based, we have proven that an online DTN routing protocol without future knowledge can perform arbitrarily far from optimal. We have also proven that optimally solving the DTN routing problem even with complete knowledge is NP-hard. Our deployment of RAPID in a DTN testbed illustrates that

our approach is realistic and effective. We have shown through trace-driven simulations using 65 days of testbed measurements that RAPID yields significant performance gains over previous work.

ACKNOWLEDGMENTS

We thank Mark Corner, John Burgess, and Brian Lynn for helping build and maintain DieselNet, Ramgopal Mettu for helping develop the NP-hardness proof, and Erik Learned-Miller and Jérémie Leguay for feedback on earlier drafts. We thank Karthik Thyagarajan for his help in formulating the Integer Linear Program. This research was supported in part by National Science Foundation awards NSF-0133055 and CNS-0519881, CNS-0721779, CNS-0845855.

REFERENCES

- [1] One laptop per child. <http://www.laptop.org>.
- [2] TIER Project, UC Berkeley. <http://tier.cs.berkeley.edu/>.
- [3] A. Balasubramanian, B. N. Levine, and A. Venkataramani. Replication Routing in DTNs: A Resource Allocation Approach. Technical Report 09-51, UMass Amherst, 2009.
- [4] N. Banerjee, M. D. Corner, and B. N. Levine. An Energy-Efficient Architecture for DTN Throwboxes. In *Proc. IEEE Infocom*, May 2007.
- [5] C. Boldrini, M. Conti, and A. Passarella. Modelling data dissemination in opportunistic networks. In *CHANTS '08: Proceedings of the third ACM workshop on Challenged networks*, pages 89–96, New York, NY, USA, 2008. ACM.
- [6] J. Burgess, G. Bissias, M. D. Corner, and B. N. Levine. Surviving Attacks on Disruption-Tolerant Networks without Authentication. In *Proc. ACM Mobihoc*, September 2007.
- [7] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE Infocom*, April 2006.
- [8] B. Burns, O. Brock, and B. N. Levine. MV Routing and Capacity Building in Disruption Tolerant Networks. In *Proc. IEEE Infocom*, pages 398–408, March 2005.
- [9] G. Casella and R. L. Berger. *Statistical Inference*. Second Edition. Duxbury, 2002.
- [10] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. In *Proc. IEEE Infocom*, May 2006.
- [11] C. Chekuri, S. Khanna, and F. B. Shepherd. An $O(\sqrt{n})$ Approximation and Integrality Gap for Disjoint Paths and Unsplittable Flow. *Theory of Computing*, 2(7):137–146, 2006.
- [12] CPLEX. <http://www.ilog.com>.
- [13] J. Davis, A. Fagg, and B. N. Levine. Wearable Computers and Packet Transport Mechanisms in Highly Partitioned Ad hoc Networks. In *Proc. IEEE ISWC*, pages 141–148, October 2001.
- [14] N. Garg, S. Sobti, J. Lai, F. Zheng, K. Li, A. Krishnamurthy, and R. Wang. Bridging the Digital Divide. *ACM Trans. on Storage*, 1(2):246–275, May 2005.
- [15] V. Guruswami, S. Khanna, R. Rajaraman, B. Shepherd, and M. Yannakakis. Near-Optimal Hardness Results and Approximation Algorithms for Edge-Disjoint Paths and Related Problems. In *Proc. ACM STOC*, pages 19–28, 1999.
- [16] B. Hull et al. CarTel: A Distributed Mobile Sensor Computing System. In *Proc. ACM SenSys*, pages 125–138, Oct. 2006.
- [17] S. Jain, M. Demmer, R. Patra, and K. Fall. Using Redundancy to Cope with Failures in a Delay Tolerant Network. In *Proc. ACM Sigcomm*, pages 109–120, August 2005.
- [18] S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *Proc. ACM Sigcomm*, pages 145–158, Aug. 2004.
- [19] E. Jones, L. Li, and P. Ward. Practical Routing in Delay-Tolerant Networks. In *Proc. ACM Chants Workshop*, pages 237–243, Aug. 2005.
- [20] J. Leguay, T. Friedman, and V. Conan. DTN Routing in a Mobility Pattern Space. In *Proc. ACM Chants Workshop*, pages 276–283, Aug. 2005.
- [21] A. Lindgren, A. Doria, and O. Schelén. Probabilistic Routing in Intermittently Connected Networks. In *Proc. SAPIR Workshop*, pages 239–254, Aug. 2004.
- [22] A. Maffei, K. Fall, and D. Chayes. Ocean Instrument Internet. In *Proc. AGU Ocean Sciences Conf.*, Feb 2006.
- [23] W. Mitchener and A. Vadhat. Epidemic Routing for Partially Connected Ad hoc Networks. Technical Report CS-2000-06, Duke Univ., 2000.
- [24] J. Ott and D. Kutscher. A Disconnection-Tolerant Transport for Drive-thru Internet Environments. In *Proc. IEEE INFOCOM*, pages 1849–1862, Mar. 2005.
- [25] C. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.
- [26] J. Partan, J. Kurose, and B. N. Levine. A Survey of Practical Issues in Underwater Networks. In *Proc. ACM WUWNet*, pages 17–24, Sept. 2006.
- [27] R. C. Shah, S. Roy, S. Jain, and W. Brunette. Data MULEs: Modeling a Three-tier Architecture for Sparse Sensor Networks. In *Proc. IEEE SNPA*, pages 30–41, May 2003.
- [28] T. Small and Z. Haas. Resource and Performance Tradeoffs in Delay-Tolerant Wireless Networks. In *Proc. ACM WDTN*, pages 260–267, Aug. 2005.
- [29] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks. In *Proc. ACM WDTN*, pages 252–259, Aug. 2005.
- [30] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Performance analysis of mobility-assisted routing. In *ACM MobiHoc*, pages 49–60, May 2006.
- [31] T. Spyropoulos and K. Psounis and C. Raghavendra. Single-copy Routing in Intermittently Connected Mobile Networks. In *IEEE SECON*, October 2004.
- [32] J. Widmer and J.-Y. Le Boudec. Network Coding for Efficient Communication in Extreme Networks. In *Proc. ACM WDTN*, pages 284–291, Aug. 2005.
- [33] Y.-C. Tseng and S.-Y. Ni and Y.-S. Chen and J.-P. Sheu. The Broadcast Storm Problem in a Mobile Ad hoc Network. *Springer Wireless Networks*, 8(2/3):153–167, 2002.
- [34] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi. Hardware Design Experiences in ZebraNet. In *Proc. ACM SenSys*, pages 227–238, Nov. 2004.
- [35] X. Zhang, G. Neglia, J. Kurose, and D. Towsley. Performance Modeling of Epidemic Routing. In *Proc. IFIP Networking*, May 2006.



Aruna Balasubramanian is a PhD student in the UMass Computer Science department since August 2005. Her research interests span Mobile wireless networks, Disruption Tolerant Networks and Energy consumption of mobile communication. Her research is partially supported by a Microsoft Research Fellowship.



Brian Levine joined the UMass Computer Science faculty in fall 1999 and is currently an Associate Professor. He received a PhD in Computer Engineering from the University of California, Santa Cruz in 1999. His research focuses on mobile networks, privacy and forensics, and the Internet, and he has published more than 60 papers on these topics. He received a CAREER award in 2002 for work in peer-to-peer networking. He has served as an associate editor of IEEE/ACM Transactions on Networking since 2005.



Arun Venkataramani has been an Assistant Professor in the Department of Computer Science at UMass Amherst since 2005. He received his PhD from the University of Texas at Austin in 2004 and was a Visiting Faculty at University of Washington before joining UMass. His interests are in networked and distributed systems and his current research focuses on mobile and wireless systems, peer-to-peer systems, Internet architecture, network security, and virtualization. He is a recipient of the NSF CAREER award and best paper awards at USENIX NSDI 2007

and 2008.