Computer Communications 32 (2009) 1699–1709

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

On the implications of routing metric staleness in delay tolerant networks

Mike P. Wittie [a,*], Khaled A. Harras [b], Kevin C. Almeroth [a], Elizabeth M. Belding [a]

[a] Department of Computer Science, University of California, 616 Mulberry Ave., Santa Barbara, USA
[b] Computer Science Department, Carnegie Mellon University, Qatar

ARTICLE INFO

Article history:
Available online 12 February 2009

Keywords:
DTN
Routing
Congestion

ABSTRACT

Delay Tolerant Network (DTN) routing addresses challenges of providing end-to-end service where end-to-end data forwarding paths may not exist. The performance of current DTN routing protocols is often limited by routing metric "staleness", i.e., routing information that becomes out-of-date or inaccurate because of long propagation delays. Our previous work, ParaNets, proposed a new opportunistic network architecture in which the data channel is augmented by a thin end-to-end control channel. The control channel is adequate for the exchange of control traffic, but not data. In this paper we present Cloud Routing, a routing solution for the ParaNets architecture. We motivate the need for such a solution, not only because of stale routing metrics, but also because of congestion that can occur in DTNs. Unable to use up-to-date routing metrics to limit congestion, existing DTN routing solutions suffer from low goodput and long data delivery delays. We show how Cloud Routing avoids congestion by smart use of forwarding opportunities based on up-to-date routing metrics. We evaluate our solution using extensive OPNET simulations. Cloud Routing extends network performance past what is currently possible and motivates a new class of globally cognizant DTN routing solutions.

1. Introduction

Delay Tolerant Network (DTN) routing protocols aim to provide the perception of full network connectivity even when an instantaneous source-to-destination path does not exist. Existing routing solutions leverage opportunistic node contacts to flood data across the network. Because flooding creates many packet copies in the network, a number of mechanisms have been developed to limit replication. These mechanisms scope flooding using hop count limits [1] or *curing* techniques [2], or exploit mobility patterns to forward data only to nodes promising more timely and more efficient delivery [3–7].

High delays inherent in opportunistic forwarding render DTN routing metrics less effective. Opportunistically disseminated information, on which routing metrics rely, can quickly become out-of-date, or stale. Worse, information from different network regions can experience varying amounts of delay and, when combined, represent a state of the network that never existed.

Existing DTN routing solutions deal with the uncertainty in routing information by replicating data. The most efficient routing solution, in terms of network resource usage, would forward only one data copy along some set of hops. Due to the uncertainty of link availability in DTNs, current solutions create many copies of the same packet so that multiple paths can be explored. However, opportunistically propagated metrics can be stale and existing solutions are forced to replicate data without considering data copies that are already close to the destination. Replication performed independently of the global traffic state creates potential for congestion as more and more superfluous copies of the same data are created in the network.

The high degree of data replication in DTNs is at odds with limited node contact time. Nodes are often in contact only for short periods of time, during which only a small amount of data can be exchanged. Under these conditions, congestion can occur when a large amount of flooded data is queued and is expected to be forwarded when nodes come into contact with each other. Because of high degree of replication, even small increases in network load result in large increases in the amount of data queued at each node. The DTN research community has often assumed that buffer space is not a limiting factor. Even so, the growth in buffered data means that more data needs to be transmitted during a node contact, increasing the probability of packets being stalled during an exchange. Stalled packets occupy network resources for a longer period of time, increasing delivery time and reducing network goodput.

The evolution of DTNs themselves may provide a solution to the fundamental problem of metric staleness and ensuing congestion. While early DTN work assumed a flat topology of mobile nodes, recent research trends have proposed more diverse architectures,

* Corresponding author. Tel.: +1 8052801690.
E-mail addresses: mwittie@gmail.com, mwittie@cs.ucsb.edu (M.P. Wittie), kharras@cs.cmu.edu (K.A. Harras), almeroth@cs.ucsb.edu (K.C. Almeroth), ebelding@cs.ucsb.edu (E.M. Belding).

0140-3664/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2009.02.006

where the opportunistic network is augmented by an end-to-end communication channel [8–10]. The possibility of end-to-end connectivity supplementing an opportunistic network creates a new space for DTN routing solutions. These new solutions could even be designed to make globally optimal routing decisions using up-to-date network state information. As a result, such solutions would successfully avoid the problems of buffer growth and poor use of node contact time.

Recent work describes DTNs that have been augmented by another network. Jain et al. suggested using multiple delay tolerant networks in unison to deliver data of different throughput and delay characteristics [11]. Surana et al. described a rural data forwarding network where monitoring information is exchanged using a Short Message Service (SMS) back channel [8]. Finally our previous work, ParaNets, considered the possibility of linking multiple networks at different layers in the OSI stack [9], or more specifically, by Chandra and Bahl at the MAC layer [12].

The core idea of ParaNets is to use multiple networks simultaneously such that each network performs only the tasks of a protocol suitable to its characteristics. We focus on the scenario where bulk data is forwarded on an opportunistic primary network and protocol control information is exchanged using an end-to-end control channel. In previous work, we have shown how a ParaNets control channel can be used to exchange data delivery acknowledgments to purge the network of superfluous copies of delivered data [9], a solution independently arrived at by Yuen and Schulzrinne [10].

In this work we explore how an end-to-end control channel can be used to eliminate metric staleness in existing DTN routing solutions. We also introduce new mechanisms that take advantage of end-to-end up-to-date network state information exchanged over the control channel to address the problems of DTN congestion and inefficient use of node contacts. We propose to reorder data queued for exchange during a contact such that packets more likely to effect deliveries are forwarded first. Our solution, *Cloud Routing*, reorders data queued for transmission using two mechanisms: *Connection Utility Reordering* and *Cloud Rank Reordering*.

Connection Utility Reordering enables efficient use of node contact time by reordering packets in order of their routing metric. While effective at lowering delay, such greedy reordering can lead to starvation of packets far from their destinations. To assure forward progress of data in the network, we add Cloud Rank Reordering.

Cloud Rank Reordering forwards data in order of its routing metric with respect to other copies of the same data residing elsewhere in the network. This mechanism enables globally aware routing decisions. Cloud Rank Reordering also assures fairness by forwarding at least one copy of each packet with the highest priority. Further, we show that using Cloud Rank Reordering and Connection Utility Reordering in combination leads to even greater network performance improvements.

While we assume the availability of a ParaNets control channel, we show that Cloud Routing performs well even when end-to-end connectivity is intermittent. This result shows we can realize the benefits of additional infrastructure and retain the generality of *disconnectivity* that characterizes DTNs.

We evaluate our work using extensive OPNET simulations of TCP-based node contact connections. The use of TCP allows us to accurately model the amount of data that can be realistically exchanged during a node contact and show the effects of DTN congestion caused by routing metric staleness. We build on these findings to demonstrate the marked benefits of employing Cloud Routing.

We compare our results against the PROPHET and Epidemic Routing solution [3,1]. While more straightforward to implement, PROPHET's performance is comparable to other contact history based routing solutions [13]. Cloud Routing achieves a factor of eight improvement in network goodput over PROPHET and a factor of two decrease in packet delay over Epidemic Routing. We also show 4720 kbps of network goodput on the ParaNets data channel for every one kbps of control channel traffic.

We conclude that the availability of an end-to-end control channel enables a fundamental shift in DTN routing and achievable performance improvements. We believe these gains make opportunistic forwarding an attractive augmentation to end-to-end networks where content dissemination over the opportunistic network is a possibility.

The remainder of this paper is organized as follows. In Section 2 we discuss current DTN routing solutions, as well as the ParaNets architecture. In Section 3, we demonstrate the challenges to DTN routing stemming from staleness of network state information. In Section 4, we detail the Cloud Routing protocol mechanisms. Section 5 describes the simulation setup used to obtain the results referenced throughout this work. In Section 6, we evaluate Cloud Routing using OPNET simulations. Finally, we conclude in Section 7.

2. Related work

A number of DTN routing solutions have been proposed by the research community. Epidemic Routing by Vahdat and Becker was the first work to achieve reliable delivery in DTNs [1]. Epidemic Routing floods data packets aggregated into bundles throughout the network. Eventually one of the bundle copies is expected to reach the destination. The work represents a proof of concept, opening the field to further study. Cerf at al. [14] standardized a more general bundle relay architecture under the auspices of the Internet Research Task Force (IRTF) Delay Tolerant Networking Research Group (DTNRG).

A different approach to DTN routing was proposed by Lindgren at al. [3]. The Probabilistic Routing Protocol using History of Encounters and Transitivity (PROPHET) maintains at each node a delivery probability vector to all other nodes. Delivery probabilities are calculated from node contact history information. If node A encounters node B in some time period, A updates its historical probability of meeting B using Eq. (1), where $P_{init} \in [0, 1]$ is an initialization constant.

$$P_{(a,b)} = P_{(a,b)old} + (1 - P_{(a,b)old}) \times P_{init} \quad (1)$$

If node A does not encounter node B in some time, node A is less likely to be a good forwarder to B and A's delivery probability should be reduced. A uses Eq. (2) to *age* its delivery probability to B with respect to the amount of time passed since last contact. Parameter $\gamma \in [0, 1)$ is the *aging constant* and k is the number of time units since last update of the delivery probability.

$$P_{(a,b)} = P_{(a,b)old} \times \gamma^k \quad (2)$$

Finally, PROPHET adjusts delivery probability using *transitivity*. If A is a good forwarder to B and B is a good forwarder to C, then A is likely to also be a good forwarder to C. Eq. (3) allows A to calculate its delivery probability to C from delivery probabilities of A to B and B to C. Parameter $\beta \in [0, 1]$ is a scaling constant that governs the impact of transitivity on a delivery probability.

$$P_{(a,c)} = P_{(a,c)old} + \beta(1 - P_{(a,c)old})P_{(a,b)}P_{(b,c)} \quad (3)$$

During contact, nodes exchange their vectors and recompute new delivery probabilities. PROPHET forwards data bundles only to nodes promising a higher delivery probability to a particular destination.

Context-Aware Routing (CAR) by Musolesi et al. uses a similar approach [4]. In addition to node contacts, CAR bases its delivery