Computer Communications 48 (2014) 44–55

Contents lists available at ScienceDirect

Computer Communications

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

Routing in hybrid Delay Tolerant Networks

Christoph P. Mayer, Oliver P. Waldhorst *

Institute of Telematics, Karlsruhe Institute of Technology, Zirkel 2, 76131 Karlsruhe, Germany

ARTICLE INFO

Article history:
Available online 4 April 2014

Keywords:
Delay Tolerant Network
Hybrid routing
Overlay Network

ABSTRACT

Delay Tolerant Networks (DTNs) have emerged as communication paradigm for providing end-to-end communication based on store-carry-forward mechanisms without the need for costly infrastructure. However, empirical studies have shown that integrating opportunistically encountered infrastructure—e.g., Internet access via WiFi—into hybrid DTNs can significantly boost routing performance. Nevertheless, extending sophisticated DTN protocols for decentralized routing towards and across the infrastructure is both complex and insufficiently understood. In this paper, we present the overlay-based *Hybrid Routing System* (*HRS*) which is—to the best of our knowledge—the first decentralized and collaborative approach for routing in hybrid DTNs that does not rely on central servers. With HRS, a large class of existing DTN protocols can benefit from opportunistic infrastructure encounters, as we show by integrating three prominent representatives of this class into HRS. In an extensive simulation study we show that (1) hybrid routing in a decentralized setting is indeed possible and can significantly boost the performance of sophisticated DTN routing protocols, (2) routing towards the infrastructure can be implemented independently for the message destination in a scalable way, and (3) communication and storage overhead can be kept low since target-oriented message forwarding across the infrastructure can avoid heavy message replication.

1. Introduction

The increasing pervasiveness of mobile devices enables a new class of networks that provide communication without costly infrastructure, solely based on the resources provided by the mobile devices and store-carry-forward mechanisms. Such *Delay Tolerant Networks* (DTN) [1–3] leverage human mobility and social behavior for providing end-to-end communication. DTNs are employed when communication infrastructure is not available, or devices are frequently disconnected such that no continuous end-to-end path may exist between sender and receiver at any time. The main challenge in DTNs is routing in face of intermittent connectivity [2]. DTN routing protocols have been proposed based on epidemic forwarding [4], social communities [5,6], or resource allocation [7]. Many current DTN routing protocols are *destination-aware* and forward messages based on (logical) *proximity* to the message's destination device, given, e.g., by probabilities of encounters [8,9], or last encounter times [10].

Mobile devices carried by humans are often multi-homed: they can participate in a DTN *and* have Internet access, e.g., via WiFi or cellular networks. This allows for integration of such infrastructure into a *hybrid DTN*. It has been shown empirically that hybrid DTNs can improve communication performance [11–13], enable communication between geographically separated DTNs, and allow for communication between devices (only) connected to the DTN and devices (only) connected to the Internet—supporting novel applications such as opportunistic computing based "Urban Sensing" [1], or "Not-so-instant Messaging" [14].

However, exploiting non-deterministic infrastructure encounters in DTN routing protocols is complex and insufficiently understood. Previous works in this area considered only very basic DTN protocols such as two-hop delivery, opportunistic flooding, and multi-copy [11–13]. Furthermore, related work did not detail on the interaction of a device connected to the infrastructure with both other devices connected at the same time or with the infrastructure itself. Often, this interaction depends on a central server [15,13], contradicting the idea of message forwarding solely based on the resources of the mobile devices. Consequently, the major challenges for using a sophisticated destination-aware protocol in a hybrid DTN is to define *collaborative and decentralized mechanisms* that can route messages (a) *towards* the infrastructure if appropriate for the message destination and (b) *across* the infrastructure to a device with higher proximity to the destination device, if available.

* Corresponding author.
E-mail addresses: mayer@tm.uka.de (C.P. Mayer), o.p.w@ieee.org, waldhorst@tm.uka.de (O.P. Waldhorst).

http://dx.doi.org/10.1016/j.comcom.2014.03.018
0140-3664/

This paper tackles these challenges and presents a methodology for destination-aware routing in hybrid DTNs, called *Hybrid Routing System* (HRS)—naturally extending the collaborative and decentralized mechanisms of existing DTN protocols into the Internet. To the best of our knowledge, this is the first routing system that allows for scalable and transparent routing in hybrid DTNs without central servers. HRS allows to bias the DTN protocol for routing towards the infrastructure as solution to challenge (a) and offers two distributed schemes that employ key-based routing overlays for forwarding messages across the infrastructure as solution for challenge (b). To allow for deployment of the DTN protocol that works best in the target scenario, HRS provides a framework for integrating a large set of destination-aware DTN protocols. To illustrate this, we integrate three state-of-the-art destination-aware protocols—*Prophet* [8], *MaxProp* [9], and *Spray*&*Focus* [10]—into HRS as proof of concept.

We conduct extensive simulation studies using real-world city maps and a model for human mobility and gain the following key findings: (1) Destination-aware routing can be effectively employed in hybrid DTNs in a decentralized setting. In fact, destination-aware DTN protocols benefit heavily from infrastructure access. For example, equipping only 30% of devices with Internet access increases the delivery probability by 169% for Prophet in the considered setting. (2) Routing towards the infrastructure can be implemented in a highly scalable way. In fact, routing decisions can be made independently of the message destination, avoiding to store proximity values for each destination currently reachable via the infrastructure. (3) Storage and communication overhead can be kept low, since heavy message replications is unnecessary due do the performance boost obtained by infrastructure availability in many cases.

This paper is structured as follows: We characterize destination-aware protocols using prominent examples in Section 2. Section 3 introduces hybrid DTNs and identifies challenges for hybrid routing. In Section 4 we present the *Hybrid Routing System* (HRS) to solve these challenges. Section 5 provides an in-depth evaluation of the performance of HRS. Related work is discussed in Section 6. Finally, concluding remarks are given.

The work presented in this paper is based on [16].

2. Destination-aware DTN protocols

Delay Tolerant Networks (DTNs) build upon a store-carry-forward paradigm using local communication only: Two devices can exchange messages when they *encounter*, i.e., when they are located in each others transmission range using a local or personal area wireless communication technology, such as IEEE 802.11 in ad hoc mode or Bluetooth. For the network model considered in the remainder of the paper we assume that several devices move in a geographic area denoted as *playground*. Devices encounter when located in each others *communication range*, which is fixed and homogeneous for all devices. Each device has an *identifier* (ID) that is fixed and unique. e.g., identifiers could comprise human-readable clear text information, like email addresses of device owners, or more abstract numbers such as *International Mobile Equipment Identity* (IMEI) as used commonly in mobile telecommunication systems. Devices can send messages to other devices—even when not located in mutual transmission range—using the device identifier to address the destination. Messages are forwarded on device encounters using the store-carry-forward mechanism.

We refer to the set of devices as $\mathcal{D}$ and use $i \in \mathcal{D}$ to refer either to the device itself, or to its ID; depending on the context. A DTN protocol delivers a message m_{ij} from a source device $i \in \mathcal{D}$ to the destination device $j \in \mathcal{D}$ using local communication and store-carry-forward. In general, DTN protocols achieve this goal with quite different mechanisms. At the one extreme some protocols use *direct delivery* that transmits a message m_{ij} only when i and j encounter directly. At the other extreme, with *epidemic routing* [4] each device $k \in \mathcal{D}$ maintains a local buffer for storing messages m_{ij}, even if it is neither the source nor the destination, i.e., $k \notin \{i,j\}$. Devices $k,l \in \mathcal{D}$ copy all messages from each others buffer upon an encounter. Obviously, direct delivery generates low overhead, while delay is high since it may take a long time until source and destination devices encounter. On the other hand, *epidemic routing* reduces delay, since it is more likely that the destination encounters a device carrying a copy of the message.[1] However, the overhead for copying and storing messages is high.

Other DTN protocols typically trade off overhead for delay by maintaining some form of routing information that helps to decide whether a particular message m_{ij} should be copied between devices $k,l \in \mathcal{D}$ upon their encounter. We categorize DTN protocols with respect to the structure and utilization of routing information as *unaware*, *self-aware*, or *destination-aware*. *Unaware* protocols do not evaluate whether device l is more likely to deliver the message to device j than devices k, but rather perform (limited) flooding as in *Epidemic Routing* [4] or *Spray*&*Wait* [17], or replication on a per-message utility as in *RAPID* [7].[2] *Self-aware* protocols evaluate the quality of devices k and l as a *forwarder in general*, irrespective of the message's destination device j. *SimBet* [18], e.g., locally manages a device-specific rating out of social similarity and betweenness, *Encounter Based Routing* [19] employs a local encounter-rate that reflects how frequently the device comes into contact with other devices, or *SANE* [20] uses social interests of device owners.

Destination-aware protocols use routing information to decide how well-suited both devices k and l are for routing the message m_{ij} *towards the destination device* j. Information required for this decision is gathered locally on device encounters, possibly taking indirect device encounters through a third device (*transitivity*) into account. Information *decays* over time (a device becomes less valuable) and is *refreshed* with further device encounters. Since the most elaborate DTN protocols can be categorized as destination-aware, e.g., [8–10], we focus on this category in the remainder of the paper.

In general form, the operation of a destination-aware DTN routing protocol can be described as follows: Each device $i \in \mathcal{D}$ maintains a *proximity* $p_i(j) \in [0,1]$ for all other devices $j \in \mathcal{D}$. It holds $p_k(j) > p_l(j)$ if device k is either more likely to encounter the destination j directly or to encounter a device m with $p_m(j) > p_l(j)$ (i.e., k will likely encounter a device m that is more likely to encounter j than l is), for $i,j,k,l,m \in \mathcal{D}$. Many proposed DTN protocols integrate the second aspect of *transitivity*. Consequently, consistent with [21] routing a message m_{ij} from device i to device j is a recursive process based on store-carry-forward. On an encounter of devices k and l, message m_{ij} carried by a device k is transferred to device l if $l = j$, or

$$p_l(j) > p_k(j). \tag{1}$$

Note that depending on the DTN routing protocol after transferring the message m_{ij} from device k to device l it can be either deleted or kept on device k. Since this leads to either a single copy or multiple copies of the message within the DTN we refer to the mode of operation of the protocol as *single-copy* or *multi-copy* mode, respectively. Single-copy mode reduces the overhead in the networks, while multi-copy mode enhances the probability of a successful delivery.[3]

[1] Assuming an ideal model of infinite device buffers and unlimited communication bandwidth.

[2] RAPID supports multiple optimization metrics, depending on the metric employed it is categorized as destination-aware or unaware.

[3] Again, assuming a model of infinite device buffers and unlimited communication bandwidth.