

Opportunistic Email Distribution and Access in Challenged Heterogeneous Environments

Tuomo Hyyryläinen, Teemu Kärkkäinen, Cheng Luo,
Valdas Jaspertas, Jouni Karvo and Jörg Ott
Helsinki University of Technology TKK, Networking Laboratory
{thyyryla, teemuk, cluo, valdas, kex, jo}@netlab.tkk.fi

ABSTRACT

The communication paradigms for delay-tolerant networking have been modeled after email. Supporting email over DTNs in a backwards compatible manner in a heterogeneous environment has yet to be fully defined. In this paper, we present a set of conventions for and extensions to the DTNRG architecture. We have implemented a system that is able to deliver emails within a DTN network, from a DTN network to the Internet, and from the Internet to the DTN network. Our system architecture includes multiple solutions for integrating traditional and DTN-based mail delivery: DTN-based messaging clients for mobile phones and PDAs, a dedicated, stand-alone DTN gateway between the Internet and the DTN network, and a personal DTN mail application proxy. The latter bridges to unmodified mail user agents running on a laptop.

Categories and Subject Descriptors

C.2.6 [Computer Communication Networks]: Internetworking;

C.2.1 [Computer Communication Networks]: Network Architecture and Design

General Terms

Design, Reliability, Experimentation

Keywords

Delay-tolerant Networking, DTN, Mobile Ad-hoc Networking, email

1. INTRODUCTION

The continuing advances in the development of mobile devices have led to many users carrying around powerful devices capable of communicating using a number of different methods such as Bluetooth and WiFi as well as storing and manipulating large amounts of data. Most of the applications designed to take advantage of these appliances rely on the device being able to form a stable, low delay connection to a server in an infrastructure network. However, the short range radio capabilities of these devices also enable

them to communicate directly with one another even when no infrastructure is available, thus forming a challenged network. With a few exceptions, traditional applications have so far failed to take advantage of the communication opportunities offered by these opportunistic contacts.

One of the most popular applications on the Internet is email. Its semantics are asynchronous and delay tolerant with users who are used to waiting hours or days for a reply—and thus a perfect fit for DTNs. However, the protocols used to transport email such as POP/IMAP and SMTP operate over TCP and therefore will not work in an environment with frequent disruptions, long delays or no end-to-end paths. Furthermore, access to the DNS infrastructure is required before messages can be transmitted.

Multiple projects have studied email over DTN, for example, in the context of tactical networks [1] and nomadic community networks [2]. The solutions in these projects have been DTN SMTP proxies that attempt to conceal the disruptions between an end user in a challenged network and a well connected mail server in an infrastructure network. While the proxy-based approach is simple to implement and deploy it fails to exploit the opportunistic communication opportunities available directly between mobile nodes. The Huggle project [3] has developed a new architecture focusing on the opportunistic side of communications with less concern about transitioning from today's Internet environment.

The IRTF DTNRG has been developing a message-based *store-carry-forward* overlay network approach [4] for enabling communications in networks with long delays and frequent disruptions. This architecture and the related *Bundle Protocol (BP)* [5] and *TCP Convergence Layer* [6] serve as a basis for our demonstration. We introduce minimal additional conventions necessary for interoperable mail transport over DTNs and for the integration with the traditional Internet.

We demonstrate email in a challenged networking scenario with mobile, heterogeneous endpoints and offer multiple solution architectures and implementations to meet different users' requirements. Instead of requiring non-disrupted infrastructure connection, we enable the nodes inside a challenged network to exchange emails even when no infrastructure is available. DTN mail gateways assist in sending and receiving mails between classical clients and servers connected to the Internet and users connected to a challenged network where an end-to-end path between these groups may never exist. The gateways also serve as a backup if direct peer-to-peer message delivery between two mobile nodes fails. Different DTN mail clients operate either as stand-alone applications or as proxies to today's unmodified Internet mail clients. Section 2 describes our system architecture and the different deployment options. The details of the implementation and the demonstration setup are presented in Section 3, followed by conclusions in Section 4.

2. ARCHITECTURE

In the traditional email architecture the end-user's mail is stored in a mailbox residing in the user's *home mail server (HMS)* in the infrastructure network. Mail destined to the user is delivered to the mailbox by SMTP and picked up by the user through POP or IMAP. This means that the user needs a stable connection to the infrastructure network in order to send or receive mail messages. Mobile nodes roaming within a challenged network might never be able to form a direct connection to the infrastructure network and therefore will not be able to send or receive email messages.

To overcome the requirement of a stable infrastructure connection, we leverage the DTN overlay network to extend the traditional email architecture to support the scenario where the end-user is connected to a challenged network. The mail messages destined to the user will be pushed into the DTN network where they will be stored and carried in the bundle caches of the DTN nodes towards the recipient. In the other direction, mail messages sent as bundles into the DTN network will be carried towards a DTN mail gateway (DTN-MGW) to be forwarded to the infrastructure network, and directly to the recipient if the recipient is also connected to the DTN network, using the same store-carry-forward mechanisms. This allows a mobile end-user to receive and send email messages while roaming in a challenged network without any infrastructure connectivity.

The general architecture for implementing the Delay-tolerant Email application is shown in Figure 1. The HMS is in practice a normal email server (shown in Figure 1 as *Mail Server*) that contains a user's mailbox serving her regular email address. The user can use POP3 to retrieve email from the HMS when connected to the Internet in the traditional mode of operation. The delay-tolerant function for HMS is implemented by using user-level filters, such as `procmail` filters. Selected incoming email messages are delivered to the DTN-MGW. The latter receives emails from trusted servers via SMTP and encapsulates them in the Bundle Protocol (BP) for forwarding towards the user with DTN connectivity.

The DTN cloud shown in shown in Figure 1 represents a challenged network composed in ad-hoc fashion out of mobile nodes. Any node within the DTN network may never have a direct connection to the infrastructure network or the DTN-MGW but the store-carry-forward mechanisms of the DTN network will enable it to exchange messages with any other node or element connected to the network.

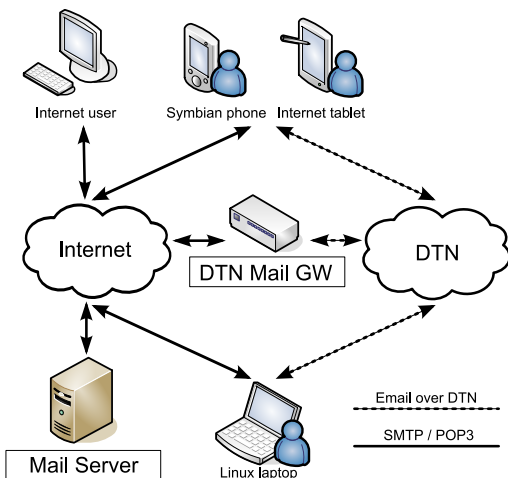


Figure 1: The architecture for delay-tolerant email

The users of the email application have thus two incoming mail routes—directly via HMS when connected to the Internet, and using the bundle protocol when having DTN connectivity. Figure 2 shows different possible configuration options we have realized for the user equipment:

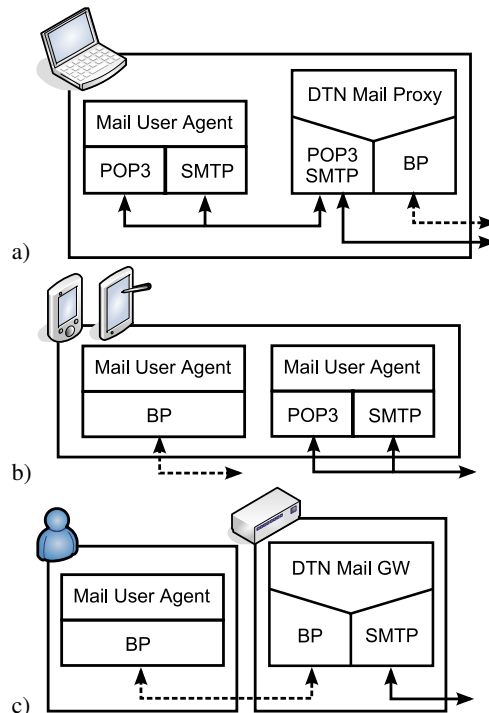


Figure 2: Options for user equipment architecture

- To enable the user to continue using traditional email applications, the user laptop can have a specialized email proxy. This proxy has both traditional and DTN connectivity options, and will receive email from both sources. It keeps track of incoming messages' message-ids so that it can filter out duplicates, and provides a POP3 and SMTP interfaces for the Mail User Agent (MUA).

The MUA can be configured to either send and retrieve all mail through the proxy or use separate mail folders and a direct connection to the mail server when connected to the Internet. For sending mail, the DTN mail proxy offers a fallback to using an Internet mail server if mail delivery via DTN has not been confirmed within a configurable time span.

- User equipment may contain separate MUAs for both connected mail and DTN mail, or
- the user equipment may contain only the DTN MUA. The latter approach requires a DTN-MGW in the network to exchange messages with Internet mail servers and is used even when directly connected to the Internet.

Beyond the configurations presented above the MUA could implement the Bundle Protocol in parallel with the traditional protocols, SMTP and POP3, with internal logic for choosing the appropriate delivery method. This, however, is a longer term solution and considering the technology transition, the approach using the

DTN-MGW provides the possibility of both traditional and DTN mail applications to coexist, without a detrimental effect on the traditional mail service.

On the user side, a separate DTN mail proxy is a transitional component, while integrating a DTN interface into the MUA is a longer term option.¹ For lightweight mobile applications, option c) is preferred. On the server side, a separate DTN-MGW can be seen as a transitional element, while integrating DTN into the mail server is a long term solution.

For sending emails over DTN, our implementations encapsulate RFC 2822 [7] compliant messages in the payload blocks of BP bundles. Mail and entity headers thus remain untouched and so does any nested structure (e.g., using MIME).

We use a simple deterministic transformation rule to convert mail addresses (which are extracted from SMTP, not from the RFC 2822 headers): The BP Endpoint Identifiers (EID) are derived from the recipient's RFC 2822 addr-spec address by prefixing it with "dtn://" and appending "/mailto" which also allows for mail-specific routing. Agents wishing to receive messages destined for a particular mail address will register the corresponding EID with their local bundle agent.

3. IMPLEMENTATION AND SETUP

We demonstrate the DTN mail application with the scenario depicted in Figure 1. As the HMS, we use our lab mail server, which is well-connected to the Internet. The server has a normal sendmail mail transfer agent.

A separate node runs the DTN-MGW software which we have developed using Java with our own JNI bindings [8] to the DTN2 reference implementation [9]. The gateway implements an SMTP client and a server for communicating with mail infrastructure on the Internet and a bundle interface for communicating with DTN nodes on the challenged network. The DTN-MGW has a wildcard registration with its local bundle router which enables it to see all the mail bundles that pass through. Routing within the DTN network is set up so that all the mail bundles get passed through the DTN-MGW's bundle router (e.g. using flooding). The DTN-MGW will convert the destination EID into an RFC 2822 addr-spec address as explained earlier and attempts to deliver the mail using SMTP. In the opposite direction, the DTN-MGW will accept SMTP connections and forward the mails as bundles. Loop prevention is done by inserting an X-DTNMail-ForwardedBy: header into the message. Any DTN-MGW will ignore a message with this header present preventing the message from being forwarded multiple times between the DTN and infrastructure networks. To prevent the DTN-MGW from being used to relay unwanted email, incoming bundles and SMTP connections should be authenticated before the mail is forwarded.

For demonstrating the heterogeneous user devices with DTN connectivity, we use one Linux laptop and one MacBook (Mac OS X)², Nokia N770 and N800 Internet tablets (Linux), and Nokia Communicators (Symbian S80).

The laptops run a DTN mail proxy and a Mozilla Thunderbird MUA. The DTN mail proxy has been developed using Java with the DTN JNI bindings. The proxy contains SMTP and POP3 services for communicating with an MUA, a local message store, and an SMTP and a POP3 client for direct mail delivery in case infrastructure connection is available.

¹We are currently developing a DTN plug-in for Mozilla Thunderbird.

²MS Windows support depends on the DTN reference implementation's availability for cygwin.

We are also using a DTN implementation [10] [11] developed in-house for Symbian based mobile phones.³ It currently runs on Nokia 9300i and 9500 Communicator devices, and has a simple user interface for receiving and sending DTN mail. The Communicators have also standard MUA, so this setup follows the configuration of figure 2b.

We have ported Berkeley DTN reference implementation to the Nokia Internet Tablet devices 770 and N800 which are running under Linux. Using the reference implementation API [9], we also implemented a native Maemo application for mail exchanging between DTN nodes. This DTN mail application is independent of the mail client shipped with the tablet (yielding, again, the configuration in Figure 2b). It directly talks with the DTN daemon by encapsulating bundle payload in MIME format.

In this setup, all communication is carried out using IP as we currently use the TCP convergence layer [6] for all DTN-based communications. For DTN message forwarding, we use static routes to the DTN-MGW (triggered whenever a device is connected to the Internet) and perform a variant of flooding among the mobile nodes to ensure maximum spread of the messages. We use BP delivery notifications to confirm message receipt and the bundle lifetime to limit the message load, peered with application-layer retransmissions and, for the DTN mail proxy, the fall-back option mentioned previously.

The above setup enables all possible variations of mail transactions: 1) incoming mail from the Internet to the challenged network through the DTN-MGW; 2) incoming mail from the Internet, with direct connection to the MUA; 3) outgoing mail from the challenged network to the Internet through the DTN-MGW; 4) outgoing mail to the Internet, with direct connection to the mail server; 5) mail transmission directly within the challenged network using BP; and 6) mail transmission from a mobile equipment to another via the mail server. Automated fall-back from DTN-based mail delivery (1, 6) to mail server-based operation (2, 4) are combined modes of operation.

Transactions 1, 3 and 5 are the interesting cases since they are not possible using the traditional mail protocols. Transactions 1 and 3 involve the passing mail among the bundle routers between the gateway and the recipient or sender of the message. This allows mobile nodes to send and receive messages to and from any node connected to the infrastructure network even if no end-to-end path ever exists between them. Transaction 5 is interesting since it allows nearby mobile nodes exchange messages even if no infrastructure connection is ever available. This is done by routing the message bundles within the DTN network directly from the sender to the recipient without the messages ever entering an infrastructure network.




Figure 3: Transaction between a traditional mail user and a DTN mail user

In transaction 1, shown in Figure 3, the mail is first sent by a

³Available at <http://www.netlab.tkk.fi/~jo/dtn/>

traditional email client to the recipient's HMS using SMTP. The HMS uses filters to forward the message to a preconfigured DTN-MGW. The DTN-MGW encapsulates the message in a bundle, derives an EID from the email address and forwards the bundle to its local bundle router from where the message enters the DTN network. Within the DTN network the message may be stored, carried and forwarded by an arbitrary number of DTN nodes. The message will finally be delivered to the recipient when the recipient's DTN enabled device makes contact with another DTN node which is carrying the message in its bundle cache. From there, depending on the type of the end-point, the message is either stored in the user's local DTN mail proxy or delivered directly to a DTN enabled MUA.

Transaction 3 is similar to transaction 1, but in the opposite direction. The DTN mail user's DTN MUA or DTN mail proxy encapsulates the email message in a bundle and forwards it to the local bundle router. The bundle is carried within the DTN network as explained earlier until it reaches a DTN-MGW. The DTN-MGW will translate the EID into an email address and attempts to deliver the message using SMTP to the recipient's HMS. The recipient will then pick up the message over POP or IMAP using a traditional email client. Figure 3 also shows transaction 3 (going from right to left), the difference to transaction 1 is that the Mail Server now represents the Internet user's HMS.




Figure 4: Transaction between two DTN mail users

In transaction 5, both the sender and recipient are connected to a challenged network over DTN and neither has a direct connection to the infrastructure network, as shown in Figure 4. Here the sender will proceed similarly to transaction 3, encapsulating and forwarding the message. However, unlike in transaction 3, the recipient is also connected to the DTN cloud allowing the network to deliver the message bundle directly. Therefore the recipient is able to receive the message without connecting to the HMS. However, as the network includes a DTN-MGW, it will also pick up the message bundle and attempts to forward it to the recipient's HMS. This way the recipient will receive the email message directly through the DTN network, but will also receive a copy in the mailbox of the HMS where it can be picked up using POP or IMAP.

In a variation of transaction 5, the sender might be connected to the DTN network but also has a stable infrastructure connection. If the sender is running a DTN mail proxy, the proxy can first attempt to transmit the message directly over DTN (in case the recipient only has a DTN connection) and fall back to SMTP delivery in case the DTN delivery fails.

As a side effect, DTN-based delivery from the Internet to the mobile node incurs *push-style delivery* which may make mail retrieval more efficient for the mobile user than the usual polling. Another side effect of the DTN-based delivery is message duplication; the recipient may receive a copy directly over DTN and later another copy from the HMS using traditional email protocols. In this case the filtering out of the duplicated must be done by the

end point MUA. The duplicate filtering requires a way to uniquely identify mail message. We use the Message-ID: header field for the identification, but other methods such as message hashes are also possible.

4. SUMMARY

This paper describes an architecture to enable mail communication in a heterogeneous environment that combines traditional server-based mail delivery and opportunistic communications for different types of devices. We have provided minimal enhancements and conventions for the DTNRG specifications and used the DTN reference implementation as well as our own implementation as a basis to create a diverse set of DTN mail clients, a DTN mail proxy and a gateway.

Our implementation environment serves as foundation for further research and development in a heterogeneous environment. We identify numerous future research topics including the disconnected (manual or automated) provisioning of `procm` forwarding rules in the DTN-MGW (e.g., to forward responses to outgoing mails); preventing spam through proper security associations between mobile nodes and the DTN-MGW (based upon S/MIME or DKIM); and, of course, proper end-to-end encryption of emails (e.g., using identity-based cryptography). On the engineering side, we are currently implementing a plug-in for the Thunderbird mail user agent to provide a convenient zero configuration environment for the user—which is important for broad acceptance across all types of devices.

5. REFERENCES

- [1] Keith Scott, "Disruption tolerant networking proxies for on-the-move tactical networks," in *IEEE MILCOM 2005*, 2005, vol. 3, pp. 3226–3231, IEEE.
- [2] Avri Doria, Maria Uden, and Durga Prasad Pandey, "Providing connectivity to the saami nomadic community," in *Proc. 2nd Int. Conf. on Open Collaborative Design for Sustainable Innovation*, 2002.
- [3] James Scott, Pan Hui, Jon Crowcroft, and Christophe Diot, "Haggle: A Networking Architecture Designed Around Mobile Users," in *IFIP WONS*, January 2006.
- [4] Vinton Cerf, Scott Burleigh, Adrian Hooke, Leigh Torgerson, Robert Durst, Keith Scott, Kevin Fall, and Howard Weiss, "Delay-tolerant networking architecture," RFC 4838, Apr. 2007.
- [5] Keith Scott and Scott Burleigh, "Bundle Protocol Specification," Internet Draft draft-irtf-dtnrg-bundle-spec-09, Work in progress, April 2007.
- [6] Mike Demmer and Jörg Ott, "TCP Convergence Layer," Internet Draft draft-demmer-irtf-tcp-clayer-00.txt, October 2006.
- [7] P. Resnick, "Internet message format," RFC 2822, Aug. 1982.
- [8] Ilkka Toje, "Dtnmail," Special Assignment, TKK Networking Laboratory, Mar. 2007.
- [9] "DTN Reference Implementation," <http://www.dtnrg.org/wiki/Code>, May 2006.
- [10] Omar Mukhtar, "Design and implementation of bundle protocol stack for delay-tolerant networking," M.S. thesis, Helsinki University of Technology, Aug. 2006.
- [11] Omar Mukhtar and Jörg Ott, "Backup and Bypass: Introducing DTN-based Ad-hoc Networking to Mobile Phones," in *Proceedings of RealMAN 2006*, May 2006.