DELAY-TOLERANT NETWORKING: MOVING TOWARDS REAL-WORLD DEPLOYMENT

Elwyn Davies
Folly Consulting Ltd
60 The Butts
Soham, Ely, Cambs CB7 5AW UK
Tel: +44 1353 624 579
e-mail: elwynd@folly.org.uk

ABSTRACT

Delay- and Disruption-Tolerant Networking (DTN) has been a research topic for more than 10 years since Vinton Cerf and NASA started working on means to expand Internet-like connectivity to situations such as communications with spacecraft on deep space missions. The distances involved and the constraints of light speed mean that the large round trip delays make conventional Internet protocols such as TCP unusable. The DTN architecture and the Bundle Protocol suite that have been developed to enable communication in challenged environments, both in space and in terrestrial situations, have now reached a level of development and testing where they are able to be used to support users in real-world situations. This paper describes some of the work that has been done to demonstrate real-world applications of DTN, especially in the N4C project in which the author was a partner, and articulates a vision of how DTN can be used to support a future Internet where information items become 'first class citizens', enabling Information Centric Networking (ICN).

1 INTRODUCTION

The conventional Internet has become nearly ubiquitous in urban areas of developed nations over the last 15 years but it has become increasingly reliant on a connected path existing between source and destination nodes when a user requests an application at the source to communicate with a destination. That is, there is an assumption of 'global reachability' in the Internet. Many of the applications that users have come to appreciate also rely on the round trip time (RTT) for data packets being quite small, so that, for example, a screen display for a web page can be built up from information retrieved by several separate requests to one or more servers or information stores. The assumption is that the upper bound on the RTT for any sort of connectivity in the Internet remains at a value where applications appear to a human user to provide near instantaneous fulfilment of requests.

There remain many areas of the terrestrial space as well as the interplanetary domain, where neither of these assumptions is valid. The term 'communications challenged' has been used for areas where the assumptions are not met, although in truth we should consider whether applying somewhat disparaging terms to the situation is appropriate. The general belief is that providing good quality communications to every area of the planet is a desirable aim, and indeed, with the way in which the operation of day-to-day life and commerce are becoming ever more reliant on efficient data communications, the digital divide will not be closed unless we can find ways to extend Internet-like communications into these so-called challenged areas.

It is clear that practical and economic realities will mean that in many areas 'challenged' networks are likely to be the only available networks. It is up to researchers and engineers to develop network architecture, hardware, and operating paradigms, together with applications and ways to use them, that allow people living and working in these areas to participate effectively in the Information Society.

In this paper we examine how the Delay- and Disruption-Tolerant Networking architecture that developed out of the NASA Interplanetary Internet initiative and described in Section 2 has been used to extend some applications into challenged areas, taking the work done by the EU Framework Programme 7 (FP7) project Networking for Communications Challenged Communities (N4C) [1] as an example. These applications are outlined in Section 3.

We will also look in Section 4 at how the networking paradigm of DTN might in future become a fundamental feature of a future Internet and especially at a vision of how DTN can support Information Centric Networking (ICN). Conclusions are drawn in Section 5.

2 NETWORK ARCHITECTURE AND APPLICATION ADAPTATION

The need for delay-tolerance in a future Internet architecture has been previously discussed [2] and the adaptation of applications to networks that do not satisfy the assumptions of the connected Internet has been considered in the N4C System Architecture [5]. Many applications do not need continuous or near instantaneous communication end-to-end. There are applications that are asynchronous such as email clients, where occasional packet exchange is the normal situation. Even apparently synchronous applications such as web browsers can often be considered to fall into the same category. At present there is a presumption that when providing communication-based services a user's *Quality of Experience* is determined by the technical *Quality of Service* in the underlying network.

Our experience in the N4C project has indicated that users who live and work in communications challenged areas have an understanding that service delivery is unlikely to be instantaneous. Accordingly they have a greater tolerance of delays and disruptions in communications than application designers for the connected Internet are normally willing to accept. Thus if applications expect that the network may not provide instantaneous connectivity and instead provide an explanation of the situation rather than treating network delay and disruption as errors that need to be fixed by the user, those users tend to report that they received a satisfactory quality of experience.

In order to support such changes in application paradigm, the network architecture needs to become delay- and disruption-tolerant. Up to now the network has mostly tried to conceal delays and disconnections that may arise due to the effects of congestion and mobility. In future, applications may need to be aware of the nature of the network to which they are attached. Application protocols may need to be redesigned removing the uniform assumption of always on, global reachability. See [5] for a taxonomy of applications varying from truly asynchronous applications (e.g., email) to bidirectional, real-time applications (e.g., Voice over IP Telephony). The latter cases will give a poor experience in a heavily delay- or disruption-prone environment but if informed about choices or adaptations of performance, the user will generally be more accepting of the situation and the services that *do* work will not be dismissed along with those that do not.

3 USING THE DTN ARCHITECTURE FOR SERVICE DELIVERY

The N4C project was an experimental research project in the EU Future Internet Research and Experimentation (FIRE) stream [6]. N4C carried out a number of experiments over three years in two test bed areas that had significant communication and power delivery challenges.

The test bed areas were

- the Kočevje region of Slovenia, and
- the Sápmi region of Arctic Sweden.

Both areas have essentially no communications infrastructure and the local environmental conditions mean that it is unlikely or impossible that permanent infrastructure will be installed in the foreseeable future.

The size of the areas, especially in Arctic Sweden meant that the only possibility of achieving communications was to rely on opportunistic encounters between pairs of nodes, at least one of which was mobile.

The DTN architecture [3] developed under the auspices of the Internet Research Task Force (IRTF) Delay-Tolerant Networking research group DTNRG [7] is designed to support, among other scenarios, data communications in situations such as those encountered in the N4C test beds. Unlike the connected Internet, data is exchanged in 'bundles' which are intended to be complete self-describing units, containing all the data needed for one side of a transaction and the meta-data needed to facilitate delivery to its intended destination. Thus delivery of an application data unit does not necessarily require multiple messages (packets) and/or multiple round trips. The basic protocol used is the 'Bundle Protocol' (BP) described in RFC 5050 [4].

The DTN architecture and the BP assume that nodes operating in a DTN environment do not necessarily have instantaneous communication with other nodes or services such as would often be assumed in the connected Internet (such as DNS servers or authentication servers), and thus have to be able to operate autonomously if necessary. Accordingly the DTN architecture addresses the architectural issues described in Section II.C of [2].

The communication paradigm can be described as 'store, carry and forward'. Bundles forwarded to a node are stored securely until another node is encountered that appears to provide a good chance of delivering the bundle to its intended destination. Then the bundle may be forwarded if the encountered node is willing. In the meantime the carrying node has a resource that may be used to fulfil information requests other than from nodes involved in the original bundle transaction.

The N4C experiments were designed to demonstrate that the DTN architecture and the BP could be used to provide real-world services in the communications challenged environments of the test beds. The services were to be integrated with equivalent services in the connected Internet allowing users operating in the challenged environment to both inter-communicate and to communicate with users attached to the connected Internet.

The primary services chosen for the N4C experiments were the basic but vitally important ones of

- Email,
- Web access, and
- Automated data collection, including web camera access.

In the event, we also provided a form of messaging service similar to the various instant messaging services available on the connected Internet. Given that messages were subject to delay, this became known as 'Not-So-Instant Messaging' (NSIM): this service proved unexpectedly popular.

The following sections describe the workings of these applications in the DTN environment in more detail. Much more detail of the experiments can be found in the deliverables of N4C at [1], especially work packages 4 and 8. There is also additional material on the N4C wiki (http://wiki.n4c.eu).

3.1 The Village Router Scenario

N4C partners Trinity College Dublin (TCD) in collaboration with Intel Labs Europe developed a system that used DTN as a transport to transfer information bundles between gateways connected to the Internet and a 'Village Router' that could be located in the area of one of the Sámi semi-permanent encampments in the Padjelanta national park in Arctic Sweden. The data could then be accessed by standard laptops or mobile phones through an access point built into the Village Router device. A simplified representation of the arrangement is shown in Figure 1.

In terms of applications, email gateways and web caching proxies were implemented on the machine at TCD, connected to the Internet, and on the Village Router, allowing access for mobile devices in the challenged area.

The various gateway machines were equipped to be Wi-Fi access points, although earlier experimentation had used *ad hoc* mode Wi-Fi but this proved to be unreliable.

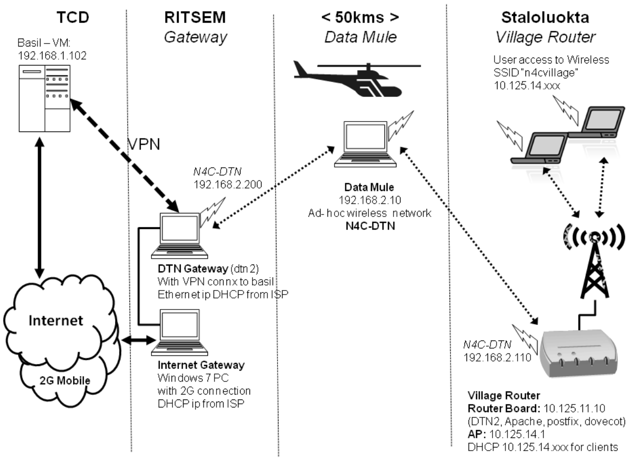

Figure 1: *Village Router Network Configuration.*

Data was mostly moved between the Internet gateway machines and the Village Routers by 'data mules'. The mules were computers with Wi-Fi clients installed in the helicopters that provide long distance transport for goods, (snail) mail, hikers, tourists and local residents during the summer period. During the summer, the alternative way of getting from place to place is walking!

The system was initially deployed with one Village Router in 2009 and a larger and longer trial was carried out in 2010 with up to five Village Routers in use at several sites and several helicopters used as data mules. The 2009 trial is described in much greater detail in [8]. A paper describing the 2010 trial is in preparation and is expected to be published in due course.

3.1.1 Village Email System

Mail servers (Postfix) were deployed on the Basil machine at TCD in Dublin and on each Village Router. The server on Basil acted as a central distribution point for mail intended for any users who had accounts on the Village Router domain. Email that arrived from the connected Internet (typically via SMTP) and from any Village Router (via DTN) was distributed to mailboxes for the users' accounts held on Basil. The mail servers on the Village Routers were effectively maintained as clones of the central mail server on Basil. Each day any changes to the mail files on Basil were parcelled up into a bundle and distributed to each of the Village Routers through DTN where it was used to update the mail files on the server.

Any mail generated at the Village Router was filed locally if it was destined for another user in the local domain and (always) sent to the central mail server, one email to a bundle.

Dovecot was used to maintain the mail files and provide an interface for the DTN mail transfer capability. The result was that users could access their email at any of the locations where a Village Router was located after a delay resulting from the time taken for bundles to be transferred via the data mules (typically one day for emails from the connected Internet).

3.1.2 Cached Web Proxy System

Caching web proxies were installed on each Village Router to act as a portal for web access from clients attached to the Village Router. A control page shown in Figure 2 was available to clients accessing the web via the Village Router.

Requests for web pages are sent as bundles to the central caching mechanism on Basil that retrieves the requested information and returns it to the portal in a bundle. The

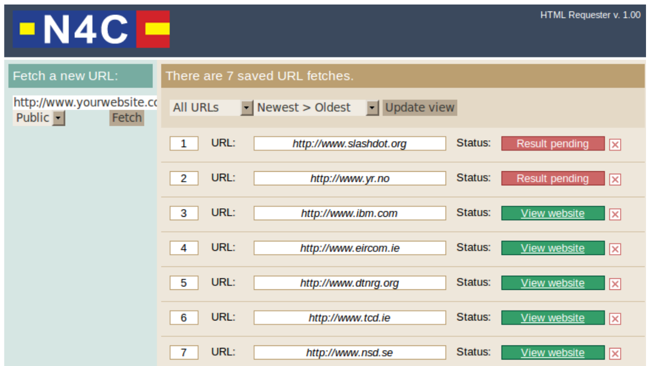

Figure 2: *Web Portal Control page. Allows users to enter a new request (top left) and monitor previous requests. Green buttons at the right indicate satisfied requests that can be retrieved from the cache. Red indicates, as yet, unsatisfied requests.*

caching mechanism attempts to get not just the single URL requested but accesses any associated URLs needed to build up the web page down to a specified search depth using 'wget' (GNU Wget tool page).

The central caching mechanism is also configured to retrieve, on a regular basis, a certain set of pages which are sent to the portals typically once per day. This is useful for supplying up-to-date weather forecasts, agricultural prices and news feeds that might be of interest to people using the portals.

There are some privacy issues to be considered for the portal. By default any requested information will be visible to any user accessing the portal, but a user can request that only the requester can see the cached results.

3.2 Nomadic Email System

As an alternative to the Village Router system, we implemented the PyMail email system, see Section 7 of [9]. This system allowed a user with a laptop computer or smartphone with Wi-Fi access to have email sent directly to the user through the DTN system. This system ran partly in tandem with the Village Router system. Individual emails addressed to nodes that had a nomadic email account were encapsulated in bundles at the central Internet gateway. The bundles were then routed towards the intended recipient machine over DTN. On arrival the email bundle was decapsulated and the bundle delivered to a local mail client through a slimline server. Outgoing email reversed the process.

The intention was that email bundles could be transported not only by the helicopter data mules but by exchange between clients who happened to encounter each other using the PRoPHET routing protocol [10].

The implementation of this process was constrained by the need to use *ad hoc* Wi-Fi during encounters between mobile nodes. It became apparent that *ad hoc* Wi-Fi implementations were mostly not ready for general deployment with some honourable exceptions. Also for all units, using *ad hoc* Wi-Fi resulted in unacceptably short battery lifetimes.

3.3 Automated Data Collection

The experiments in Slovenia organized by N4C partners MEIS ran nearly continuously for three years (and are still in progress). The main focus of these experiments was collecting meteorological and other environmental data from a network of sensors in the Kočevje region of Slovenia.

The collecting stations were autonomous and entirely automatic. Data was collected on a regular basis by a 'drive past' from a DTN equipped car and returned to another DTN node at MEIS' office where the data was recorded and analysed.

In addition to the drive past system, MEIS also developed a passive DTN node, the hardware for which is just a standard USB memory 'key'. When plugged into a suitably configured (static) computer, bundles were exchanged between the DTN node on the static machine and the node on the key. Thus the key could be used as a general purpose data transporter or mule that was cheap and entirely standard hardware.

Such nodes could be despatched to the central point by (snail) mail. This system was used for a radiation monitoring project linked to the worries of some local residents about pollution of their homes from a nearby industrial plant.

3.4 Outcome of Experiments from the User Point of View

The round trip time in the N4C networking experiments was typically 1 to 2 days. For the data collection experiments data was delivered on average half a day after collection. In general terms the residents in the challenged areas who used the system were well-used to information arriving more slowly than typical urban dwellers and were happy with the experience.. The main problem is finding an economically viable way of providing an ongoing service.

Clearly there are additional services (notably e-commerce and banking) which it would be helpful to support, but these services will require additional work to provide secure operation at the application layer. Existing Internet security protocols require continuous connectivity and are not applicable to DTN scenarios.

4 FUTURE WORK

The applications used in the N4C experiments are reliant on application specific gateways to translate between packets and bundles. If the whole internet is to become delay-and disruption-tolerant, it would be desirable for applications to work directly with bundles. The bundles could then traverse the network end-to-end and the IP network would become a transport substrate for the parts of the DTN network that were continuously connected. Interestingly, this sort of scheme has presaged the transformation of the communications infrastructure at the last major changes of the network (e.g., dial-up networking over POTS to ADSL).

It may also help if applications can know and communicate to servers that they are operating in a DTN region. This might, for example, allow web servers to provide an alternative mode of service that avoids multiple round trips and the effects of program generated content.

The 'store' aspect of the 'store, carry, and forward' DTN paradigm offers a way of implementing the ideas of Information Centric Networking (ICN) [11]. ICN envisages data items becoming first class citizens either besides or instead of nodes. A DTN network already supports the sort of cached, addressable data items envisaged by ICN.

5 CONCLUSIONS

DTN as envisaged by RFC 4838 [3] has now progressed to the stage where it is usable for real-world deployment of applications. The N4C experiments have shown that DTN segments can be integrated into the connected Internet and allow at least some communication services to be provided reasonably seamlessly across the boundary.

The need to service communications challenged regions does not appear to be one that will disappear over time, so that DTN could prove to be an architectural technique that will be increasingly relevant in the future.

References

[1] N4C Consortium, Networking for Communications Challenged Communities web site. http://www.n4c.eu 2008-11.

[2] J. Ott. Delay Tolerance and the Future Internet. *The 11th International Symposium on Wireless Personal Multimedia Communications (WPMC'08),* Lapland, Finland, 2008.

[3] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H.Weiss. Delay-Tolerant Network Architecture. RFC 4838, April 2007.

[4] K. Scott and S. Burleigh. Bundle Protocol Specification. RFC 5050, November 2007.

[5] E. Davies. N4C System Architecture. *Deliverable 2.1, N4C Consortium.* Folly Consulting Ltd, 2011.

[6] European Commission. Future Internet Research and Experimentation (FIRE) web site. http://cordis.europa.eu/fp7/ict/fire/.

[7] DTNRG members. IRTF Delay-Tolerant Networking Research Group (DTNRG) wiki and web site. http://www.dtnrg.org 2004-2011.

[8] S. Farrell *et al.* Report on a Arctic Summer DTN Trial. *Journal Wireless Networks*, Volume 17 Issue 5, July 2011.

[9] K. J. Grøttum, S. Sjursen, J. Näslund, E. Davies, and B. Grašič WP3 Evaluation Report. *Deliverable 3.2, N4C Consortium*, February 2010.

[10] E.Davies (ed) *et al.* Probabilistic Routing Protocol for Intermittently Connected Networks (PRoPHET). draft-irtf-dtnrg-prophet-09, April 2011.

[11] D. Kutscher, B. Ahlgren, H. Karl, B. Ohlman, S. Oueslati, and I. Solis. Dagstuhl seminar Proceeding 10492 – Information-centric networking. *Dagstuhl Seminar on Information-Centric Networking*, Dagstuhl, Germany, September 2010.