Sharing the dream

The consensual networking hallucination offered by the Bundle Protocol

Lloyd Wood, Peter Holliday and Daniel Floreani
London, England;
Brisbane and Adelaide, Australia
lloyd.wood@yahoo.co.uk,
pho@tpg.com.au, daniel.floreani@internode.on.net

Wesley M. Eddy
Federal Network Systems
Verizon Business
Cleveland, Ohio, United States
wesley.eddy@verizonbusiness.com

***Abstract*—We reconsider desirability of the Bundle Protocol (BP) as a universal solution for Delay- and Disruption-Tolerant Networking (DTN). The BP is intended to provide a single solution that is applicable to a wide variety of differently-challenged DTN networks, even though those networks are unlikely to interact with one another. This paper asks whether such a single protocol can encompass all varied DTN networking needs. It asks whether attempting to repeat the previous success of the homogeneous Internet by layering over all networks is suitable for the heterogeneous DTN world, where diverse application needs and operational requirements lead to diverse, scenario-specific, applications. This position paper is intended to encourage discussion of the role, scope, and adoption of the BP.**

Keywords-Bundle Protocol, Delay-Tolerant Networking, DTN.

I. INTRODUCTION

Delay- and Disruption-Tolerant Networking (DTN) began with work on the *Interplanetary Internet* – an attempt to move space agencies towards packetized networking, rather than streaming data using low-level link protocols. The long propagation delays of interplanetary distances led to the name *Delay-Tolerant Networking* [1]. Designing protocols for use at long distances requires careful attention to how timing and timers are handled; protocols must be patient, and proactive. The predictability of space probes, knowing when the Earth will be next visible, can take advantage of advance scheduling.

The DTN concept was extended by Fall to include challenged networks in general [2]. This included military ad-hoc networking scenarios, which attracted DARPA attention as *Disruption-Tolerant Networking*. Here link connectivity and availability can be far less predictable than for deep space, of shorter propagation delay, and on more varying channels. This is a very different environment, suggesting a very different networking approach. Neither the Interplanetary Internet nor military ad-hoc resembles the terrestrial Internet – one thing these scenarios do have in common. A shared approach is that off-the-shelf Internet protocols are not usable, and radical new approaches are required. Both scenarios have advocates – the Consultative Committee for Space Data Systems (CCSDS) for defining ISO standards for deep space, and DARPA for military work. This has led to considerable funding and effort in those two communities to support a solution which can work across both scenarios – even though the scenarios are so very different [Fig. 1]. The Bundle Protocol (BP) was originally intended for deep space, and is now considered *the* single solution for all DTN scenarios by the IRTF DTN Research Group, which develops the BP. But is the BP suitable for this wide-ranging role across all, very varied, DTN networks?

II. THE BUNDLE PROTOCOL AND USE IN DTNS

The BP exists in a simple layered architecture, talking to the application above, and passing bundles to and from a network-specific *convergence layer* below [Fig. 2]. Nodes supporting the BP may rely on a number of convergence layers suitable for different networks and network conditions. The BP ignores all properties of the network it travels across; that is the convergence layer's problem, as is mapping BP-specific bundle addresses and routing, itself still under research, onto the locally-understood network. The BP has no protocol timers other than a bundle expiry time, allowing it to work with the long delays of deep space. The BP was not designed to include reliability checks of either its own protocol headers or payload data. Reliability mechanisms such as error-checking and retransmission are problems left to the convergence layer and applications outside the BP itself, which is not concerned with these issues. The BP can be described as a complex extensible container format, with optionally secured payloads, carried by the supporting local network infrastructure. Networking features not yet provided by the BP are listed elsewhere [3].

Just as the Internet suite provided a way for fixed computers with different types of wired links to interact, the BP suite is intended to provide a way for everything computational to interact over every possible link type, wired or wireless.

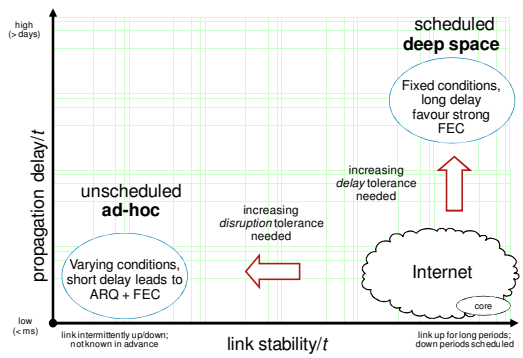

Figure 1. How military ad-hoc and deep-space scenarios differ

9781-4244-3941-6/09/$25.00 ©2009 IEEE

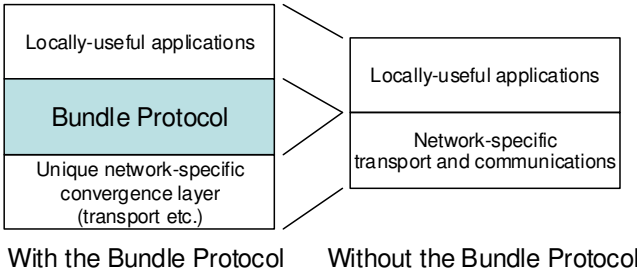

Figure 2. Layering with and without the Bundle Protocol

The Internet Protocol (IP) overlaid a variety of link and subnet technologies, by providing a shared address space that allowed wired (and now wireless) computers to interact in interesting and useful ways and share applications. (Telnet and ftp were the original 'killer apps' before the web browser). The BP is intended to overlay links and networks, including the Internet, so that all the devices using the BP can interact and share applications, or at least share a common software abstraction. The BP provides a new shared *consensual hallucination* (coined by William Gibson for 'cyberspace') for thinking about networking in scenarios not considered by the original Internet inventors, just as the Internet Protocol provided the consensual hallucination and common overlay for computers to share resources, files, and applications.

As well as the civil space and military efforts discussed earlier, the DTN BP is being trialed in *e.g.* networks to support nomadic reindeer herders [4]. Use of a common BP could allow a deep space node to communicate directly with a reindeer herder – but what do these DTN networks carry that is useful across these very different operational scenarios? What can be said about reindeer to a Mars rover? DTN applications are often scenario-specific to meet diverse needs, and their data is not useful to other DTNs. Diversity of operational needs, networks and applications leads to diverse convergence layers, routing protocols, and so on, leading to incompatibility.

DTN nodes providing Internet-like functionality to remote users will need to interact with the Internet. Many DTN sensor networks and nodes will also need to interact with and provide information to nodes on the terrestrial Internet. Much collected sensor data requires large amounts of post-processing before it becomes useful information; this computation is often done in terrestrial Internet networks, *e.g.* for remote sensing data [5].

But many separate DTN networks with different uses, particularly sensor networks, are often private, isolated networks with custom technologies implemented in each DTN node. Those networks do not need to interact with one another, and their applications will be focused on their local tasks. Including the BP adds relatively little utility to such networks.

The BP provides commonality with a shim between application and convergence layer – but this commonality is not essential locally, nor is it needed for communication with existing Internet nodes [Fig. 2]. This commonality is only needed when talking directly between very different DTN networks that share the BP abstraction, across very different operational scenarios running different local applications.

III. COST/BENEFIT TRADEOFFS

Implementing the BP in any DTN adds development effort. There are clear use cases for having a DTN network interact with the existing Internet, *e.g.* a sensor sending data to and receiving instructions from Internet computers. However, the utility of DTN nodes in one network interacting directly with DTN nodes on another network, with different operational scenarios that have different purposes, but requiring the common BP to communicate, is unclear. It is reasonable for multiple space probes, built by different manufacturers, to interact amongst themselves and use the BP, just as different ad-hoc nodes in a military deployment can do – but these two civil space and terrestrial military domains are unrelated. The Interplanetary Internet (ironically, planned with the BP in place of IP) has little in common with the needs of the US Department of Defense; a Mars probe has nothing to say to a tank commander on Earth. The scenarios and domains do not interact. CCSDS has begun adding to the BP itself, to meet its own needs. It is reasonable to expect that DARPA will do likewise. What will remain common across these bespoke BPs?

Reusing select appropriate Internet technologies in DTN networks has benefits: decreasing new development effort and cost and improving interaction with the Internet. This has been shown by daily use of IP for delivery of remote–sensing data from space. The BP was also tested in this scenario, but added little to the already disruption-tolerant IP-based operations [5]. By reusing existing Internet technologies, DTN networks in different operational scenarios can readily connect with the Internet, and interoperate with everyday Internet applications. They can also interact with each other through the Internet *if* it is desirable to do so, and *if* application data can be shared.

What will be the universal 'killer app,' the compelling new functionality akin to telnet or ftp, which will encourage widespread adoption of the BP across different scenarios?

It is possible to extend the consensual hallucination of the Internet to new frontiers and to reuse appropriate Internet technologies, *e.g.* UDP-based transport protocols or MIME file transfers, to enable DTN networks to interact directly with the Internet. An entirely new consensus, where different scenarios and applications adopt the BP and use it to interact directly between themselves, seems unwarranted. The dream of a grand unifying common BP, with interoperation between different DTN scenarios, requires that the BP shows clear operational or implementation benefits for all local scenarios. That dream is being explored by researchers. Will their dream be realized?

REFERENCES

[1] V. Cerf *et al.*, "Delay Tolerant Network architecture: the evolving Interplanetary Internet," expired internet-draft, August 2002.

[2] K. Fall, "A Delay-Tolerant Network architecture for challenged Internets," SIGCOMM, pp. 27-34, Karlsruhe, Germany, August 2003.

[3] L. Wood, W. M. Eddy and P. Holliday, "A Bundle of Problems," IEEE Aerospace conference, Montana, March 2009.

[4] A. Lindgren and A. Doria, "Experiences from deploying a real-life DTN system," 4th Consumer Communications and Networking Conference (CCNC), 2007.

[5] W. Ivancic *et al.*, "Experience with delay-tolerant networking from orbit," accepted by the International Journal of Satellite Communications and Networking for publication, 2009.