Ad Hoc Networks 23 (2014) 80–86

Contents lists available at ScienceDirect

Ad Hoc Networks

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

Space for Internet and Internet for space

Scott Burleigh [a,*], Vinton G. Cerf [b], Jon Crowcroft [c], Vassilis Tsaoussidis [d]

[a] *Jet Propulsion Laboratory, California Institute of Technology, United States*
[b] *Google, United States*
[c] *Computer Lab, University of Cambridge, United Kingdom*
[d] *Space Internetworking Center, Democritus University, Greece*

ARTICLE INFO

Article history:
Received 15 May 2014
Accepted 21 June 2014
Available online 3 July 2014

Keywords:
Space
Internet
DTN

ABSTRACT

Space flight and Internet service are technologies that are currently complementary but seem to be on the verge of integration into a new "space internetworking" discipline. The authors believe a comprehensive realization of space internetworking technology could dramatically enhance space exploration, augment terrestrial industry and commerce, benefit the economically disadvantaged, and nurture human and civil rights.

1. Introduction

People think of space flight and the Internet as distinct technologies, both exciting, both even central to the unfolding of our future as a species over the centuries ahead, but fundamentally independent. It is the thesis of this paper that this independence is ending. An infrastructure for universal interconnection among humans and their artifacts is emerging which will unify internetworking and space flight technologies, enhancing the impact of both on the way we enjoy and use our home in the solar system.

2. Space flight operations to support Internet

Internetworking is made possible by "links" that copy digital information between computers' memories. These links can take the form of copper telephone wires, television service cables, fiber optic lines, or, more recently, radio signals among cellular telephone towers. But all such terrestrial link technologies are innately limited by geography. Links that are formed by radiation to and from orbiting satellites, however, can transcend all geographic barriers with equal ease.

The potential value of such links has been recognized for decades. In the early 1980s an experimental satellite connection between ARPANET and European hosts was established, operating at 64 kbps and confirming that Van Jacobson's congestion control algorithm for TCP could tolerate .72 s of signal propagation latency [1,2].

By now, of course, satellite-based Internet service is a multi-billion-dollar industry, despite technological obstacles of its own. However, it can be argued that we have only begun to explore the ways in which spacecraft can enhance the capabilities of the Internet.

In particular, we continue to use satellites only as analogs for wired infrastructure, conducting continuous end-to-end TCP/IP dialogues. An alternative would be to exploit the enormous potential bandwidth of satellites in low-Earth orbit functioning as "data mules", acquiring data from possibly disconnected customers over high-frequency radio or free-space optical links, storing the data in local memory, and physically transporting the data over orbital tracks until they overfly gateways into wired local area networks – or the wired Internet – that can forward the data over high-speed terrestrial links.

* Corresponding author. Tel.: +1 818 393 3353; fax: +1 818 393 6871.
E-mail address: Scott.C.Burleigh@jpl.nasa.gov (S. Burleigh).

http://dx.doi.org/10.1016/j.adhoc.2014.06.005
1570-8705/

This model is a variation on well-known "sneakernet" concepts that have bemused network designers for many years. For example, we have long been advised to "never underestimate the bandwidth of a 747 filled with DVDs." The theoretical capacity of a Boeing 747 filled with Blu-ray disks flying from New York to Los Angeles is 245.8 terabits per second [3]. The latency of this transmission channel – a round-trip time of around 12 h – would make it unsuitable for many Internet applications, but not all: remote system backups and the delivery of online journals, for example, do not depend on sub-second round trip times.

Less fancifully, we consider a constellation of hundreds of low-Earth orbiting (LEO) satellites with free-space optical links to the ground stations in their field of view, each operating at 1 Gbps. This constellation would constitute a geography-independent communication channel whose bandwidth would be on the order hundreds of gigabits per second. Again the round-trip latency would be high, but less so than for the network of 747s: LEO satellites travel at 25–30 times the speed of a commercial jet aircraft, so the theoretical round-trip time from New York to Los Angeles would be on the order of half an hour. At that rate, even Facebook posts and movie downloads might be suitable applications.

Here, then, we note that network technology to automate communications over unusual network topologies like those discussed above – where no end-to-end continuous connectivity is available at any time – already exists. Delay-Tolerant Networking (DTN [4]) protocols readily handle these kinds of topologies, simply because at the time the protocols were designed they needed to do so in order to enable generalized internetworking in space flight operations, as discussed in the next section.

3. Internet to support space flight operations

The operations of spacecraft, whether crewed or robotic, have always been reliant on communication with mission teams on Earth. Early spacecraft had little or no on-board computing power, and mission communications were limited to relatively simple signaling; communications could be managed manually. But over the past 20 years the on-board computational capability of spacecraft has increased rapidly. With that increase has come rapid growth in the power and complexity of flight software, finally enabling collaborative operation among multiple spacecraft in situ. Moreover, observational instrument payloads on spacecraft are now able to generate science information at very high rates. Taken together, these advances have resulted in rapid growth in the volume and sophistication of mission communications and in the potential complexity of the mission configurations those communications support [5].

Just as in terrestrial research communications in the mid-20th century, the increasing demands on flight mission communication infrastructure are making earlier methods of communication management increasingly untenable. It is quickly becoming important to automate the operation of flight mission communication channels by deploying network technology.

However, the specific technology solutions that addressed the terrestrial research communication problem and developed into the Internet we know today are in some ways unsuitable for internetworking in space flight missions.

Interplanetary communication is characterized by long and variable signal propagation latencies (e.g., 1¼ s to the Moon; from 4 to 20 min to Mars) and by frequent lengthy lapses in connectivity, caused by the interposition of a planetary body between source and destination and/or by transient considerations of power or attitude management that make reception impossible. Consequently:

- Protocols must be connectionless. In the length of time required for all the round trips required to establish a connection, the communication opportunity might terminate.
- Timer expiration intervals are not predictable from statistics: at any time a lapse in connectivity might add minutes or hours to the round-trip time.
- Since lapses in connectivity are routine and nominal, they must not be interpreted as changes in network topology that need to be propagated to routing tables.
- In the worst case, there may never be contemporaneous connectivity among all nodes on the end-to-end path from source to destination. End-to-end forwarding and retransmission as performed by TCP/IP would never succeed.

In short, the operational assumptions on which the design of Internet transport, network, and routing protocols are based do not hold.

It is these considerations that led to development of the DTN architecture to which we alluded above. DTN protocols solve these problems by accepting different design assumptions:

- The DTN "Bundle Protocol" forwards data much as the Internet Protocol does, except that outbound bundles are held in local storage until forward links are available, destinations are expressed as names rather than addresses (because the topological location of the destination node might change while data are en route to it, so it may be necessary to delay the binding of the destination to a specific address), and routing decisions are based not on knowledge of current network topology – which is, in the general case, unavailable – but on expectations of future network topology.
- The retransmission of lost data is performed within the network rather than from end to end. Having survived the perilous transit from one node to the next, a DTN "bundle" need never repeat that ordeal: if it is lost on the next leg of its journey it is retransmitted from the last node at which it was successfully received.
- DTN security focuses on protecting bundles not only while they are in flight but also while they are at rest, because they may so often and so long be at rest. Security measures are integrated directly into the bundle structure rather than imposed only in ephemeral encapsulating structures.