Experience with delay-tolerant networking from orbit

Will Ivancic
Satellite Networks and Architectures Branch
NASA Glenn Research Center
Cleveland, Ohio, United States
wivancic@grc.nasa.gov

Lloyd Wood and Peter Holliday
Global Government Solutions Group
Cisco Systems
London, United Kingdom and Brisbane, Australia
lwood@cisco.com, phollida@cisco.com

Wesley M. Eddy and Dave Stewart
Verizon Federal Network Systems
at NASA Glenn Research Center
Cleveland, Ohio, United States
weddy@grc.nasa.gov, dstewart@grc.nasa.gov

Chris Jackson and James Northam
Spacecraft operations
Surrey Satellite Technology Ltd (SSTL)
Guildford, United Kingdom
C.Jackson@sstl.co.uk, J.Northam@sstl.co.uk

***Abstract*— The Disaster Monitoring Constellation (DMC), constructed by Surrey Satellite Technology Ltd (SSTL), is a multi-satellite Earth-imaging low-Earth-orbit sensor network where captured image swaths are stored onboard each satellite and later downloaded from each satellite payload to a ground station. The DMC is currently unique in its adoption of the Internet Protocol (IP) for its imaging payloads and for satellite command and control, based around reuse of commercial networking and link protocols. Earth images are downloaded from the satellites using a custom IP-based high-speed transfer protocol developed by SSTL, *Saratoga*, which works well in unusual link environments. Store-and-forward of images with capture and later download during passes over ground stations gives each satellite the characteristics of a node in a Delay/Disruption Tolerant Network (DTN). DTNs are under investigation in an Internet Research Task Force (IRTF) DTN research group (RG), which has developed a 'bundle' architecture and protocol. We experiment with use of this DTNRG bundle concepts onboard the UK-DMC satellite, by examining how *Saratoga* can be used as a convergence layer to carry the DTN Bundle Protocol, so that sensor images can be delivered to ground stations and beyond as bundles. This is the first use of the Bundle Protocol from orbit. We use our practical experience to examine the strengths and weaknesses of the Bundle architecture for DTN use, paying attention to fragmentation, custody transfer, and reliability issues. We similarly examine and discuss an alternative network stack, based around proposed use of the Hypertext Transfer Protocol (HTTP) that we have been architecting, which we believe has potential applications across a range of DTN networks. We use our practical experience to make suggestions about DTN use and adoption of IP in sensor networks.**

Keywords- *DTN, delay-tolerant networking, Saratoga, bundle, space, low Earth orbit, LEO, reliability, timestamps.*

I. INTRODUCTION

Delay/Disruption Tolerant Networking (DTN) has been defined as an end-to-end store-and-forward architecture capable of providing communications in highly-stressed network environments. A "bundle" protocol (BP) sits at the application layer of some number of constituent internets, forming a store-and-forward overlay network and service [1]. Key capabilities of the BP include:

- Custody-based retransmission - the ability to take responsibility for a bundle reaching its final destination.
- Ability to cope with intermittent connectivity.
- Ability to cope with long propagation delays.
- Ability to take advantage of scheduled, predicted, and opportunistic connectivity (in addition to continuous connectivity).
- Late binding of overlay network endpoint identifiers to constituent internet addresses [2].

The DTN protocol suite is intended to consist of a group of well-defined protocols that, when combined, enable a well-understood method of performing store and forward communications. DTN can be thought of as operating across varying conditions across several different axes, depending on the design of the subnet being traversed:

- low or high propagation delay.
- dedicated or shared, congested links.
- links with intermittent disruption and outages or scheduled planned links.

In a low-propagation-delay environment, such as may occur in near-planetary or terrestrial environments, DTN bundle agents can utilize chatty underlying Internet protocols, such as TCP, that negotiate connectivity and handshake connections in real-time. In high-propagation-delay environments such as deep space, DTN bundle agents must use other methods, such as some form of scheduling, to set up connectivity between the two bundle agents, and can use less chatty transfer protocols over IP.

Low Earth Orbit (LEO) is a low-propagation-delay environment of less than ten milliseconds delay to ground, with long periods of disconnection between passes over ground stations. For the UK-DMC satellite, contact times consist of 5 to 14 minutes per pass with one or two available ground station contact times per 100 minute orbit – assuming multiple available ground stations. The ground stations are connected across the terrestrial Internet, which has different operating conditions (congestion-sensitive, always on) from the private links between each satellite and ground station (intermittent but scheduled, and dedicated to a single download traffic flow).

II. THE RATE MISMATCH PROBLEM

Fig. 1 illustrates a LEO satellite ground network with a DTN Bundle Agent sink located at a remote location. The final remote location for the downloaded imagery could be a satellite control station and office or a laptop 'in the field' with wireless connectivity – it really doesn't matter.

In this example, an image is to be transferred from the DTN source, the LEO satellite, to the DTN sink. In this example, the hypothetical image file is too large to be transferred during one pass over a single ground station. Rather, three passes are required to transfer the complete file to ground. These passes could all be via the same ground station or could utilize three different ground stations. The minimum time a complete image file could be transferred using a single ground station is a little over 300 minutes, assuming one pass per 100-minute orbit. However, using three different ground stations, the entire image could be downloaded in a fraction of an orbital period, by downloading fragments of the image to each ground station and reassembling the complete image file on the ground.

If some type of rate-based file transfer is used between the sink and source, problems will arise if ground link capacity does not match or exceed the rate of the space-to-ground link; the transfer becomes limited by any bottleneck in the path. In order to increase the download rates across each link, the transfer can be split into multiple separate hops, where the download is stored and forwarded locally across each hop – note, this is the situation whether using a single ground station or multiple ground stations.

The requirement is to get the image off the spacecraft as efficiently as possible, as spacecraft pass time is the major constraint, and then transfer separately across the different environment of the terrestrial Internet afterwards. The DTN BP is one example of a protocol that enables such functionality with its separate convergence layers, and can thus compensate for rate mismatches between the private wireless link and the shared path between ground station and remote destination.

III. UK-DMC SATELLITE CHARACTERISTICS

The UK-DMC satellite is one of five similar imaging satellites currently launched into low Earth orbit in similar sun-synchronous planes. It was launched in September 2003, with a design lifetime of five years. This imaging constellation continues to grow, with at least four more satellites to be added in the next two years to maintain a continuous on-orbit imaging capability. While these satellites are government-owned, the UK-DMC satellite is also used to provide imagery for commercial resale when not otherwise tasked in imaging campaigns or supporting disaster relief. Anyone may request an image and pay the associated costs [3].

The UK-DMC satellite is primarily operational rather than experimental. However, SSTL has also run experiments onboard the UK-DMC such as investigating GPS reflectometry

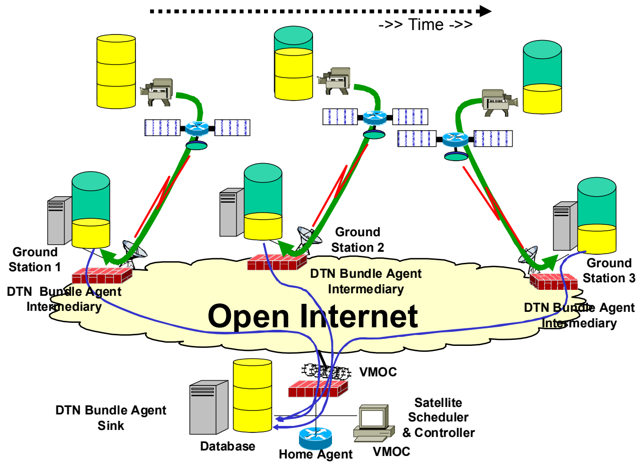

Figure 1. DTN Ground Network

[4,5] and networking experiments have taken advantage of an onboard Internet router [6,7]. SSTL continues to permit NASA Glenn Research Center (GRC) to use the UK-DMC satellite for experimentation with new forms of networking.

The UK-DMC satellite's onboard payloads include:

- The Cisco router in Low Earth Orbit (CLEO). CLEO has been used for network testing and is its own experiment to simply show that a commercial-off-the-shelf router could survive and function in orbit. CLEO is not used for DTN bundle testing.
- Three Solid-State Data Recorders (SSDRs). One SSDR, based around a StrongARM processor, supports the onboard GPS reflectometry experiment. Two SSDRs of a newer design, with Motorola MPC8260 PowerPC processors, support the imaging cameras. One of these SSDRs is used for DTN testing. These run the RTEMS operating system, which supports the POSIX API and BSD sockets. These have a constrained operating system firmware size limit of 1 MByte, and storage capacities of 1 GByte and 512MByte RAM respectively.

There is a downlink of 8.134 Mbps, and a command uplink of 9600 bits per second – this is highly asymmetric. Both links use the standard IPv4/Frame Relay/HDLC encapsulation developed for space by Keith Hogie [8]. IPv6 has also been tested over these links, using the onboard CLEO router [9,10].

The IP-based transport protocol used for downloading images is SSTL's original implementation of *Saratoga*, retroactively called version 0, running over UDP, the User Datagram Protocol. *Saratoga* version 0 is the existing operational SSTL file transport protocol, originally developed to replace and improve transfer performance rates over an implementation of CCSDS CFDP that was previously used by SSTL. *Saratoga* version 1 is an improved specification, with enhancements to *Saratoga* version 0, which has now been documented publicly as a contribution to the IETF [11].

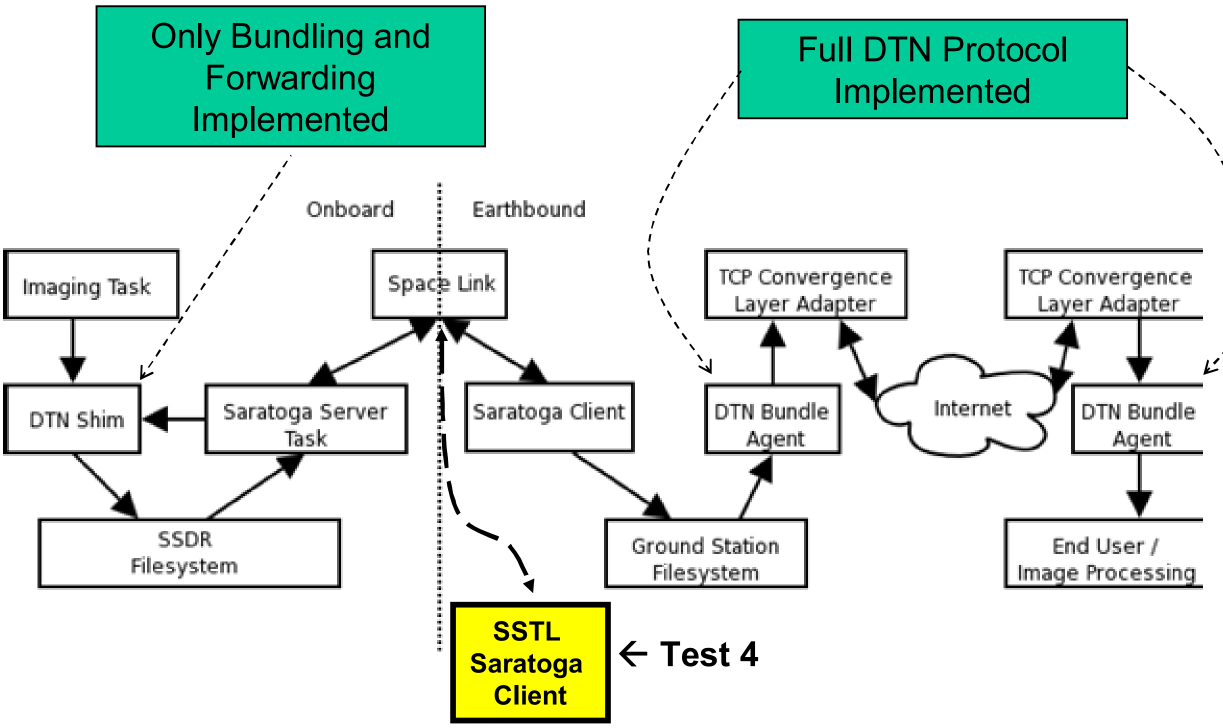

Figure 2. UK-DMC Implementation

Saratoga is designed to fill a private link at maximum available link rate. It uses efficient selective negative acknowledgements to minimize the amount of return traffic needed to ensure reliable delivery. A single TCP (Transmission Control Protocol) flow cannot fill a private link due to TCP's assumptions about congestion, slow start, and back off. Once the forward channel/acknowledgement backchannel capacity ratio exceeds 50:1, TCP's performance is degraded. *Saratoga* continues to perform at link rate for a forward/back capacity ratio on the UK-DMC of over 850:1. Later DMC satellites use faster downlinks and higher link asymmetry ratios.

IV. Experimental Bundling Implementation

A. For the UK-DMC Satellite

Fig. 2 illustrates how DTN bundling is implemented onboard the UK-DMC satellite and in the ground infrastructure. Support was added to *Saratoga* to allow it to act as a bundle transport convergence layer on the space-ground link [12]. Only the bundle-forwarding portion of DTN was implemented onboard as a simple networking "shim" since available code space is constrained, and a goal was to have the onboard DTN implementation be transparent to normal UK-DMC operations, living side-by-side with the existing operational code in a non-disruptive manner. This was considered acceptable for testing as the UK-DMC acts only as a source of DTN data, and does not need to receive and parse bundles from elsewhere.

Thus, the DTN-bundle-receiving intelligence only needed to be present in the ground station implementation of the *Saratoga* client and the DTN bundle agent. The *Saratoga* client in the ground station queries the UK-DMC satellite for a directory of files, and then requests any files with a ".dtn" extension and an associated satellite image file. The satellite image file and associated metadata files are transferred to the ground, where the *Saratoga* client reassembles the bundles and then presents them to the full DTN bundle agent - full DTN-2 bundle agent implementations were used both at the ground station and the final DTN destination [13]. Finally, to demonstrate proactive fragmentation, the DTN fragments were reassembled at the final DTN destination.

B. Ground Development and Testing

A schematic diagram of the DTN ground testbed, where bundling over *Saratoga* was prototyped, is given in Fig. 3. This development testbed, which reused the CLEO ground-based testbed duplicating in-orbit UK-DMC hardware, requires:

- The PowerPC-based Solid-State Data Recorder (SSDR) that resides in the Cisco router in Low Earth Orbit (CLEO) engineering model, where the bundle file is generated.
- A channel emulator that emulates the 9600 bps uplink and the 8.134 Mbps downlink. This uses a Spirent SX-14 data link simulator to provide channel delay and bit-error-rate emulation independently on both the uplink and downlink.
- A DTN bundle agent acting as the ground station. This bundle agent queries the DTN source onboard the SSDR for files and bundles using the SSTL *Saratoga* version 0 file transport protocol.

- A remote bundle agent sink for DTN bundles.

All network layer communications used IPv4, with the simulated space/ground data link implemented using standard Frame Relay/HDLC encapsulation. This enables easy interfacing to Cisco 3640 routers.

C. *Overall Goals of These DTN Experiments*

The goals of the experiments were to:

- Demonstrate that DTN code and general SSTL code can coexist without affecting normal SSTL spacecraft or ground station operations;
- Demonstrate DTN bundle transfers from UK-DMC to SSTL; and,
- Demonstrate proactive fragmentation of DTN bundles.

The ability to run DTN bundling without affecting normal SSTL operations would enable the DTN bundling code to remain loaded as part of the operational system, so that NASA would not need to take the UK-DMC satellite away from normal operations for dedicated experimental use. Lack of impact on normal imaging operations results in significant cost savings for future tests and demonstrations.

V. TESTS

A. *Test Configuration*

We determined that an image size of roughly 160 Mbytes would allow us to run a full 160-Mbyte file transfer, a 160-Mbyte DTN bundle transfer, and two 80-Mbyte DTN bundle fragment transfers during a single satellite pass over a ground station. Demonstrating normal DTN bundle transfers verifies DTN operation and shows that *Saratoga* can be used as a bundle convergence layer. Proactive fragmentation is required to perform large file transfers over multiple ground stations.

SSTL instructed the UK-DMC satellite to acquire a 150-Mbyte image while passing over the Gulf of Khambhat, India on 25 January 2008 at ~04:35 UTC.

B. *Bundles on the UK-DMC Satellite*

Fig. 4 shows how bundles were created onboard the UK-DMC satellite. When the image of the Gulf was acquired, the large 150-Mbyte image was stored in the SSDR and automatically named by the operating system.

The SSTL naming convention is to use a 10 character name for the recorded image. Here, the name was DU000c76pm. As well as this file, the DTN shim created three additional files of approximately 70 to 80 bytes. These files are the DTN bundle headers containing the DTN metadata.

The first DTN bundle header contained metadata for the entire 150-Mbyte file. The second two DTN files contained DTN metadata used for proactive fragmentation. The arbitrary convention used to name the metadata files was to use the default system name with an extension of ".dtn" added to the full bundle name. For the fragmentation dtn metadata files, the system default name along with the start and stop file offset and the ".dtn" were used. For the 150-Mbyte satellite image, this resulted in the two proactive fragment metadata files shown in Fig. 5. Only a very small amount of additional metadata and filespace was required to perform proactive fragmentation.

C. *Results of DTN Tests*

Three UK-DMC satellite passes were taken to test the latest NASA/Cisco/SSTL firmware code supporting *Saratoga*/DTN bundling.

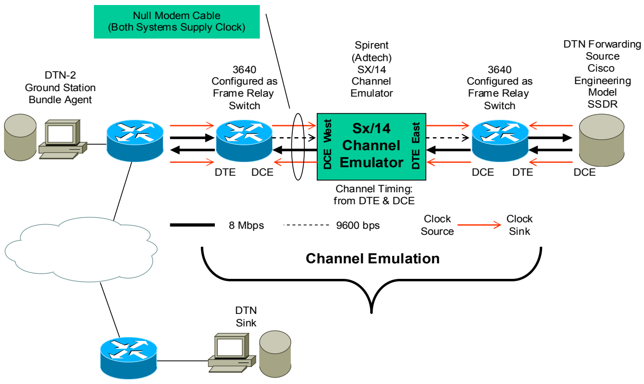

Figure 3. DTN Testbed

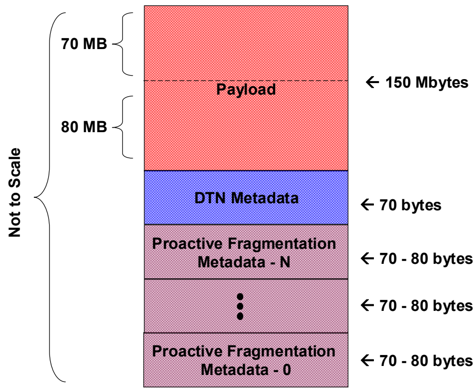

Figure 4. Bundles on the UK-DMC

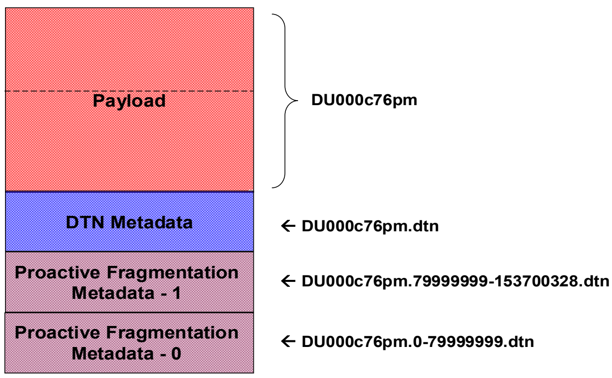

Figure 5. File Naming Convention

The passes occurred as follows on 25 January 2008:

- 07:54 - 08:07 UTC
  28 degrees maximum elevation.
- 09:31- 09:45 UTC
  45 degrees maximum elevation.
- 11:12 - 11:21 UTC
  5 degrees maximum elevation.

Four tests were performed:

- Basic image file download, using existing *Saratoga* file transfer techniques (NASA GRC's implementation of *Saratoga* version 0).
- Download of the same file as a DTN bundle.
- Download of the same file, using DTN proactive fragmentation with 80-Mbyte preconfigured fragments
- Normal file transfer using SSTL's workstation and SSTL's implementation of *Saratoga* version 0. This provided an operational control to be compared with the first three experiments [Fig. 2].

For test 1, the satellite image file, DU00076pm, was received at the SSTL ground station in Guildford, England using NASA Glenn's implementation of *Saratoga* version 0. This file was then transferred to NASA GRC over the public Internet using normal file transport protocol (FTP).

For test 2, the satellite image file, DU00076pm, and associated DTN metadata file for the full bundle, DU00076pm.dtn, were received by the *Saratoga* client on the ground and presented as a full bundle to the bundling agent, Bundling-SSTL, at SSTL's ground station. This was resent as a full bundle to the NASA Glenn Research Center DTN sink, Bundling-GRC1.

For test 3, proactive fragmentation, the first proactively-fragmented bundle file from the UK-DMC was received on the ground by the *Saratoga* client, the fragmentation bundle was reconstituted and presented to the DTN bundle agent, Bundling-SSTL. This bundle fragment was then automatically transferred from Bundling-SSTL to Bundling-GRC1 using DTN. The second proactive fragmentation bundle was not retrieved. Upon further investigation, the directory and the *syslog* file onboard the UK-DMC indicated that the first fragmentation metadata file was created, but not the second. Post-experiment analysis showed that SSTL's operating system limits file names to 32 characters. This is a settable parameter. The file name, DU000c76pm.79999999-153700328.dtn, is 33 characters long and thus the file was not created.

Initial results showed all image files reconstructed at the GRC DTN bundle sink had the correct file size, but that the file contents did not match, as there were long strings of zeros in various places in each file. The placement of these long strings of zeros differed for each file.

SSTL performed an additional control test, test 4, where the ground station computer running the GRC bundle agent and *Saratoga* client was replaced by one with SSTL's normal *Saratoga* client [Fig. 2]. That copy of the 150-Mbyte image was downloaded without errors.

On the first pass, tests 1 and 2 were successful regarding operation of DTN and the ability to either use either *Saratoga* for straight file transfers or *Saratoga* with bundling to transfer DTN bundles between the UK-DMC payloads and the ground, demonstrating bundle delivery from space. Also, the DTN-2 forwarding agent, Bundling-SSTL, was able to automatically forward the DTN bundles to a DTN-2 bundling agent at NASA Glenn Research Center, Bundling-GRC1. It was then possible to then extract the image file from the DTN bundle.

The post-test analysis revealed a number of minor problems in the experiments conducted. The reconstructed DTN bundle payload and image file (tests 1 and 2) did not match. The DTN bundling and forwarding worked, but there was a problem in the NASA GRC implementation of the *Saratoga* client regarding filling holes in missed data. Retransmission requests were not performed properly. This programming problem has since been found and fixed. A programming problem was also found in the DTN-2 code implementation put on the SSTL bundling agent in the ground station, as a bundle became stuck in a temporary file and was never transferred to GRC.

VI. Known Problems and Issues

A. *Reliability, Error Detection, and Checksums*

The current Bundle Protocol specification is recognised not to address reliability, in that it has no checksum support for error detection and rejection of corrupted bundle payloads or bundle metadata [14]. That means that one cannot determine if the bundle information received at each hop was received error-free, or has been stored safely for long periods.

Error detection is a very basic networking concept that was overlooked in the bundle protocol design, which ignores the well-known end-to-end principle. The current proposed workaround is to use the bundle security specification and to wrap the bundle using a reliability-only cipher rather than a security cipher that provides a reliability check as a side-effect of security [15].

The bundle security specification was not implemented here. Thus, there was no support for reliability checks. Had checksums been part of the core DTN bundle specification, the "holes to fill" implementation problem would have been discovered early on, and corrupted bundles would not have been transferred through our entire DTN network.

B. *Time Synchronization Problems*

During initial ground testing it became clear that network time synchronization is critical for DTN, which assumes that all communicating DTN nodes understand local UTC time. This is probably not a reasonable requirement for many DTN networks, as most DTN networks will be nondeterministic. Furthermore, DTN is a network overlay at the application layer that may be running on top of ad-hoc networks in highly stressed environments. The requirement to synchronize all DTN node clocks is not necessarily practical or deployable. However, in this scenario, with scheduled LEO passes over a ground station, it is necessary for everything to know what the time is to support the pass opportunity. The question is: how much clock drift should be permissible?

This clock synchronization problem was experienced during initial ground testing. All DTN bundle agents were originally configured and tested at NASA GRC in Cleveland, Ohio. One bundle agent was sent to Guildford, England. A second was sent to Universal Space Network (USN) in Alaska.

When performing initial DTN bundle transfers from SSTL to GRC to USN, it was noted that the machine clocks had drifted sufficiently enough to result in the bundle time stamps being out of synchronization. The DTN bundles were therefore rejected due to time-stamp mismatch. Once the machines were resynchronized, DTN transfers operated correctly. This problem was articulated at the 71st Internet Engineering Task Force meeting in March of 2008. Others have noted similar problems [16]. If a remote sensor in a DTN network loses its clock for any reason, it cannot communicate with other DTN nodes using the bundle protocol, as they will reject its query bundles requesting the right time. Some other, additional, method is thus required for bootstrapping correct clock timing to re-enable bundling on that sensor node.

VII. Other Approaches to DTN Networking

Other approaches to delay-tolerant networking do not require the bundle protocol.

One simple approach, leveraging existing standards, is to use HTTP as a transport-layer-independent 'session layer' between two communicating DTN nodes, hop-by hop [17]. New *Content-Source* and *Content-Destination* headers are added, which provide routing information end-to-end. As HTTP servers must reject transfers with unknown Content-headers, this creates a separate DTN network that will not affect existing traditional web use of HTTP. Reuse and implementation of HTTP in this way appears straightforward.

Fitting HTTP to *Saratoga* for long-delay or private networks is possible; HTTP can already be used over TCP across the shared, congested, Internet. HTTP provides the ability to easily transfer content identified by MIME [18]. Such content identification is something that is also missing from the existing bundle protocol. It can be argued that the web and email have become successful partly because it is easy to determine what application should be invoked to receive a delivered file, due to their universal adoption of MIME.

VIII. Conclusions

Delay-tolerant networking bundle transfers have been demonstrated from orbit. The DTN bundling shim onboard the UK-DMC and the ground station *Saratoga* client and bundle reconstitution mechanisms should continue to operate without affecting normal UK-DMC operations, giving NASA access to an operational DTN and bundle testbed in orbit.

Some minor software implementation bugs regarding retransmission of errored packets and file name limitations were identified and have since been corrected and ground tested. We are awaiting further on-orbit testing opportunities.

The lack of end-to-end reliability checksums in the Bundle Protocol and its need for clock synchronization have been shown to be real deployment issues during our initial tests. We hope that these architectural issues will be examined in future versions of the DTN architecture and bundling specifications.

The DMC satellites and their use of the Internet Protocol for imaging transfers provide working operational examples of how IP can be used effectively for sensor networks, to allow easy integration with the terrestrial Internet for data delivery. This use of IP to carry sensor data performs well on a daily basis, without requiring the Bundle Protocol.

IX. References

[1] K. Scott and S. Burleigh, “Bundle Protocol Specification,” IETF RFC5050, November 2007.

[2] V. Cerf *et al.*, “Delay-Tolerant Network Architecture,” IETF RFC 4838, April 2007.

[3] DMC International Imaging: http://www.dmcii.com/

[4] M. Unwin, S. Gleason and M. Brennan, “Sensing Ocean, Ice and Land Reflected Signals from Space: Results from the UK-DMC GPS Reflectometry Experiment,” ION GNSS 18th International Technical Meeting of the Satellite Division, 13-16 September 2005, Long Beach, California.

[5] S. Gleason *et al.*, “Processing of bistatically reflected GPS signals from low Earth orbit for the purpose of ocean remote sensing,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 43, No. 6, pp. 1229-1241, June 2005.

[6] W. Ivancic *et al.*, “Secure, Network-Centric Operations of a Space-Based Asset: Cisco Router in Low-Earth Orbit (CLEO) and Virtual Mission Operations Center (VMOC),” NASA Technical Memorandum 2005-213556, May 2005.

[7] L. Wood *et al*. “Using Internet nodes and routers onboard satellites,” International Journal of Satellite Communications and Networking, volume 25 issue 2, pp. 195-216, March/April 2007.

[8] K. Hogie, E. Criscuolo and R. Parise, “Using standard Internet Protocols and applications in space,” Computer Networks, special issue on Interplanetary Internet, vol. 47 no. 5, pp. 603-650, April 2005.

[9] W. Ivancic *et al.*, “IPv6 and IPSec Tests of a Space-Based Asset, the Cisco router in Low Earth Orbit (CLEO),” NASA Technical Memorandum 2008-215203, May 2008.

[10] L. Wood *et al.*, “IPv6 and IPsec on a satellite in space,” conference paper B2.6.06, 58th International Astronautical Congress, Hyderabad, India, September 2007.

[11] L. Wood *et al.*, ”Saratoga: A Scalable File Transfer Protocol,” work in progress as an internet-draft, draft-wood-tsvwg-saratoga, February 2008.

[12] L. Wood *et al.*, “ Using Saratoga with a Bundle Agent as a Convergence Layer for Delay-Tolerant Networking,” work in progress as an internet-draft, draft-wood-dtnrg-saratoga-03, February 2008.

[13] DTN reference implementation, October 2007 release, available from http://www.dtnrg.org/wiki/Code

[14] K. Fall and S. Farrell, “DTN: an architectural retrospective,” IEEE Journal on Selected Areas in Communications, vol. 26 no. 5, pp. 828-836, May 2008.

[15] W. M. Eddy, L. Wood and W. Ivancic, “Checksum Ciphersuites for the Bundle Protocol,” work in progress as an internet-draft, draft-irtf-dtnrg-bundle-checksum, March 2008.

[16] W. M. Eddy, “DTN Time Sync Issues,” email to the IRTF dtn-interest mailing list, 1 April 2008, and subsequent discussion.

[17] L. Wood and P. Holliday, “Using HTTP for delivery in Delay/Disruption-Tolerant Networks,” work in progress as an internet-draft, draft-wood-dtnrg-http-dtn-delivery, February 2008.

[18] N. Freed and N. Borenstein, “Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies,” IETF RFC 2045, November 1996.