THE INTERNATIONAL SPACE STATION, OPTICAL COMMUNICATIONS, AND DELAY TOLERANT NETWORKING: TOWARDS A SOLAR SYSTEM INTERNET ARCHITECTURE

Alan Hylton[1], David Israel[1], Miriam Wennersten[1], Rachid Chaoua[1], Mikael Palsson[2]
Mark Sinkiat[3], Greg Menke[3], Daniel Raible[4], Rachel Dudukovich[4], Nadia Kortas[4]
Brian Tomko[4], Ethan Schweinsberg[4], Prash Choksi[4], Jose Lombay-Gonzalez[4], Tom Basciano[5]
Bill Pohlchuck[6], Josh Deaton[7], Jennifer Sager[8], Mahima Kaushik[9], Tad Kollar[10]

alan.g.hylton@nasa.gov

[1]NASA Goddard, Greenbelt, MD, USA

[2]Omitron Inc., Greenbelt, MD, USA

[3]Space Coast Aerospace Services LLC., Cocoa, FL, USA

[4]NASA Glenn, Cleveland, OH, USA

[5]NASA Johnson, Houston, TX, USA

[6]Boeing Company, Houston, TX, USA

[7]NASA Marshall, Huntsville, AL, USA

[8]KBR Wyle, Greenbelt, MD, USA

[9]Columbus Technologies and Services Inc., Greenbelt, MD, USA

[10]Banner Quality Management Inc., Fairview Park, OH, USA

Keywords: DTN, ILLUMA-T, Cloud computing

Abstract

As Delay Tolerant Networking (DTN) matures as a software product its use cases have extended to infrastructural and architectural studies, bringing DTN closer to a widespread, operational technology. While it is known that a DTN can be configured, given complete information about a system of nodes, to provision a system with a complete network capability, what remains is to tackle practical considerations such as scalability, networking best practices, and how to establish service providers. In this paper, we document progress towards a scalable DTN architecture that was tested across multiple organizational and project boundaries; our approach was tested across nodes flying on the International Space Station (ISS) connected to a ground network.

The network consists of four main network areas: the ISS, the White Sands Complex (WSC), the Mission Cloud Platform (MCP), and NASA's Huntsville Operations Support Center (HOSC). The connections on the ground (WSC, MCP, and HOSC) are straightforward technically, but represent administrative and non-technical challenges that had to be overcome. The connection from the ISS to WSC was realized through a hybrid optical/RF means. In particular, NASA's Integrated Laser Communications Relay Demonstration (LCRD) can form a relay from low-Earth-orbit (LEO) to the Earth, and was used in conjunction with the LCRD LEO User Modem and Amplifier Terminal (ILLUMA-T) on the ISS to form a link to one of three geographically diverse ground stations. A major goal was to achieve communications while keeping the ISS nodes unaware of ground station choices and scheduling.

These disparate links were networked using High-rate DTN (HDTN), which enabled the desired architecture. The network architecture, tests, HDTN performance metrics, and general observations made are all discussed in detail, including a discussion on the successful utilization of the DTN security standard known as DTN Bundle Protocol Security (BPSec). We conclude with suggestions for next steps, and in particular focus on extending this architecture to the Near Space Network (NSN) and the upcoming LunaNet.

1 Introduction

Delay Tolerant Networking (DTN) refers to a suite of protocols (e.g. [1][2][3]) designed to bring networked communications to space communications. The goal is the Solar System Internet (SSI). The standards and their implementations are necessary ingredients - and come with their own requirements - but on their own do not give rise to a network. Indeed, the way such a network is built and operated is of vital importance, yet largely remains open. As the software and underlying protocols mature we can begin the study of the systems that use these building blocks. An experiment campaign using the International

Space Station (ISS), space-borne optical communications, and a multi-faceted ground component was conducted using DTN as the unifying agent to advance DTN network architectures. This paper discusses the experiments, the results, and next steps.

The current plans for the upcoming lunar internet, dubbed LunaNet, features a multi-hop multi-path network with multiple space network service providers[4]. The experiments detailed in this paper prove an architectural approach designed with space network service provision in mind, and can now serve as a platform to extend towards more complex systems.

1.1 Networking

Networking enables scalability, and in fact this is the primary purpose. The importance and urgency of scalable communications in space is accelerating, and successful protocols that enable the growing numbers must also support a growing heterogeneity of nodes[5].

The diversity of an SSI has many dimensions. The computers running DTN may have different architectures, performance characteristics, and operating systems. The links might be radio, optical, or fiber/copper. The links also might have one-way light times ranging from milliseconds to minutes (or more). At a higher level, some nodes might be leaf nodes, some might be routers. Taking the router point to the next step, different nodes might have different owners and administrators.

An SSI must work across multiple project, programmatic, and governmental boundaries.

1.2 Physical Layer

In DTN, we must consider both the machines running the DTN software (as well as the DTN software itself) and the links connecting the nodes. For this test campaign, several types were used.

1.2.1 Space links: The Laser Communications Relay Demonstration (LCRD) is a currently-operating NASA mission in geosynchronous orbit that among other capabilities provides optical and radio frequency (RF) communications to ground stations and optical communications to assets in Low Earth Orbit (LEO). LCRD can operate above 1Gbps and has a round-trip time (RTT) of roughly 4 seconds[6][7].

LCRD was used to connect the DTN component on the ISS to the ground segment using the Integrated Laser Communications Relay Demonstration Low-Earth Orbit User Modem and Amplifier Terminal (ILLUMA-T) payload, which was integrated into the International Space Station (ISS) and is now defunct[8]. ILLUMA-T connects to the onboard Ethernet network and LCRD, providing a 1Gbps Ethernet connection to the ground stations [9].

1.2.2 Terrestrial links: Connections on the ground used Ethernet and Internet.

1.2.3 Aerial links: An airplane was used as an end-point, for many of the experiments. This aircraft was piloted around Lake Erie in Ohio, and connected via laser to an optical ground station at the NASA Glenn Research Center (GRC). The system is similar to [10].

1.3 Paper Layout

Before the experiments are laid out, brief discussions of the system, DTN, and the goals are given in Sections 2, 3, and 4 respectively. The architecture is then discussed at length in Section 5. The tests and results are given in Section 6. We conclude with future steps in Section 7.

2 The System

A non-trivial network was constructed across five separate domains; refer to Figure 1.

2.1 Space Component

Aboard the ISS two nodes were flown. There was a leaf node using the Telescience Resource Kit (TReK), a library of software for space systems which includes DTN**trek**. There was also a gateway node which was used to form a bridge to the ground segment.

2.2 LCRD Network

Neither ILLUMA-T nor LCRD run DTN; rather these form a link from the ISS to the ground. More precisely, while connected to ILLUMA-T, LCRD can establish a connection to one of the three ground stations (two optical, one RF). The LCRD network has four DTN nodes - one per ground station and a gateway node to connect to other ground networks.

2.3 Huntsville Operations Support Center

The Huntsville Operations Support Center (HOSC), which traditionally manages all US communications to and from the ISS hosted two DTN nodes for this experiment. One was a gateway node, and the other a leaf node using TReK.

2.4 DTN Engineering Network

The DTN Engineering Network (DEN) is a proving ground for large-scale DTN tests hosted at NASA GRC. A dedicated network of four nodes was stood up for these tests, and included a gateway and a node on an airplane.

2.5 Cloud Gateway

With an increasing emphasis on cloud computing, the gateway between the LCRD network, the HOSC, and the DEN was a DTN node running in the Amazon Web Services (AWS) Govcloud.

3 Delay Tolerant Networking

Introductions to DTN abound, and because our emphasis is on the network itself we will not attempt to give details here beyond the bare minimum. The interested reader is referred to [11].

3.1 Overview

DTN is an overlay network. This means that there can be *any* network between nodes that at the DTN layer are adjacent; terrestrial nodes might connect to each other over the Internet (and in our case some do). Connections over space links might use protocols designed for space, such as the Licklider Transmission Protocol (LTP) [12], which is a reliable (and heavily tunable) protocol. These intermediary protocols are known as convergence layers (CLs).

3.2 Data

The primary unit of data in a DTN is the *bundle*, which is used for end-to-end transport. A bundle can effectively be of any size, and is depending on the particular characteristics of the CLs used the bundles will need to be made compatible with them via convergence layer adapters (CLAs). For variety, we may refer to a DTN node as a bundle agent. The modern version of the bundle protocol (BP) is BPv7[2], though BPv6[13] is still used.

3.3 Addressing

Scalability in networks using the Internet Protocol (IP) benefits for the hierarchical structure of IP addresses. In DTN, addressing remains an active area of study (see e.g. [14]), and flat naming is used instead. DTN names used in this experiment uses numerical (integer) names. These names are referred to as Endpoint Identifiers (EIDs).

3.4 Routing

There are multiple routing algorithms designed for DTNs, but contact graph routing (CGR) is the most common; see e.g. [15]. CGR depends on contact tables that detail when any two nodes can communicate (and in which direction). Based on these data, end-to-end routes are computed using a modified Dijkstra's algorithm. This means that when a bundle is sourced or received, CGR will be invoked to compute a path to the destination and based on this path the next bundle hop is identified.

3.5 Security

The specification for security, DTN Bundle Protocol Security (BPSec), provides confidentiality and integrity. Confidentiality is achieved by a cipher-suite and is end-to-end (as opposed to hop-by-hop). Integrity is also performed end-to-end, and ensures that any modification (malicious or otherwise) to the bundle payload can be detected. See [16] and [3] for details.

3.6 Reliability

DTN offers a high-layer reliability mechanism called *custody transfer*, which can be thought of as hop-by-hop bundle-layer acknowledgements. The details are many and intricate, and in fact custody is only defined for BPv6.

4 Goals

In 2013, DTN was demonstrated over the Lunar Laser Communications Demonstration (LLCD). This was a step forward for DTN as well as laser communications in space, as DTN can low-pass filter out link disruptions and anomalies; however, all bundling was performed on Earth. These tests remain important as they allow us to recognize how DTN glues together a system of systems, but to give follow-on experiments greater impact (such as ours) it was decreed that there was to be "no cheating," meaning that bundling must occur *everywhere* and in a multi-hop multi-path network.

On top of this, the architectural component - our primary focus - must be emphasized. We now consider an operational reality of the ISS: even updating contact schedules for an experimental DTN can take weeks of review. But recall that LCRD might connect the ISS to one of three ground stations, each of which has a unique name. A contact schedule would be very complex (read: hard to manage) if all individual contacts were listed. Making this worse, handovers (say in the event of weather) would not be possible. This means that the architecture must somehow manage these realities in a way that is transparent to the ISS.

These observations led to the following:

1. Bundling must occur in space as well as on the ground.

   Indeed, DTN must operate in all environments in order to connect them.

2. The network must be comprised of multiple subnetworks that are separated by project and/or programmatic boundaries.

   Without network subdivision, the architecture cannot yield scalability.

3. We must contain the complexity of LCRD.

   For ILLUMA-T and LCRD to provide meaningful communications to the ISS, it must be done in a way that requires as few reconfigurations as possible to the ISS while offering all ground stations and the ability to perform ground station handovers as needed for continuity of service.

4. The experiments must give rise to future interface control documents (ICD), concepts of operations

(CONOPS), and formal requirements.

> These experiments must enable the next evolutionary steps of DTN and the SSI built on DTN, which will require these elements.

As noted for the ISS, any requirement of its nodes to have awareness of LCRD's portion indicates an unsustainable and unextensible architecture. Looking ahead to LunaNet and beyond, we expect the lunar network to have multiple service providers spread across multiple countries. Therefore, all of these goals lead to the notion of space network service provision.

5 Architecture

The entire network (at the DTN layer) is depicted in Figure 1. For the purposes of this paper many of the details are distractions, so we restrict ourselves to a high-level overview. The TReK nodes (2 and 31) use the Interplanetary Overlay Network (ION) DTN implementation. Node 30 uses DTN Marshall Enterprise (DTNME). All others use High-rate DTN (HDTN). The dashed lines represent connections to either optical ground stations (OGS) or the RF ground station known as the Payload to Ground Link Terminal (PGLT), and are the only links to use LTP. All other links use the Transmission Control Protocol CL (TCPCL).

In the contact schedule the links are *always up*. HDTN allows for in-band methods to determine if a link is active or not using LTP pings; this method allowed the ISS Gateway Node (Node 1) to determine in real-time if it could send or receive bundles from the LCRD network; this alleviates much operator burden and in fact enables a shift in operational practices by encouraging users to send data regardless of link state as DTN can automatically store, carry, and forward the bundles as appropriate.

Traditionally, a contact graph for CGR means globally-distributed and globally-consistent information, meaning every link schedule is known by every node in the network. Despite scalability concerns, from a CGR perspective (the aforementioned ISS review process notwithstanding) the topology in Figure 1 can absolutely be implemented this way. This would, however, prove nothing and be a menace to maintain and modify. Returning to the ISS reconfiguration woes, this means an alteration in say the DEN would induce week-long headaches.

Putting the ISS configuration on hold, we show the first steps towards scalability. For our tests, the TReK nodes only saw each other as sources and destinations; all other nodes were either routers or inconsequential. Hence it is unnecessary for, say, Node 31 to maintain the link schedule for the DEN. Figure 2 depicts Node 31's contact plan:

1. Any bundle sourced at Node 31 *must* get routed through Node 30, so Node 30 must be in Node 31's plan.
2. Likewise, any bundle received by Node 31 *must* have come through Node 30.
3. Node 31 needs to know that Node 30 (the only next hop) can find a path to Node 2, but does not need to know the details.

These observations lead to a greatly simplified contact plan that is unique to Node 31. Note that any intermediary change (say, to Node 20) has no impact on Node 31's plan. This approach to routing tables is useful for scalability (particularly in terms of maintenance) but is not itself architectural or even novel in nature, and rather serves as a warm-up[1]. Now, we return to the ISS Gateway Node, Node 1. For emphasis, we recall that any configuration changes on the ISS follow lengthy processes. Counter-intuitively, this is actually quite liberating; indeed, this induces the following in the contact plan:

1. The contact should target one EID.
2. The contact should *always* be on.

The "always-on" component is straightforward - HDTN allows a link it expects to be established to be probed with pings; this was the course taken. The singular EID could be achieved by giving all ground station DTN nodes the same EID, or only having one ground station DTN node and routing all IP traffic to it. However, both of these approaches violate the "no-cheating" clause.

Giving all ground station nodes the same EID leads to several issues:

1. It precludes operating with multiple users simultaneously.
2. Internal routing is impossible (say, during a handover).

[1] It also leads to a lesson-learned; see Section 7.

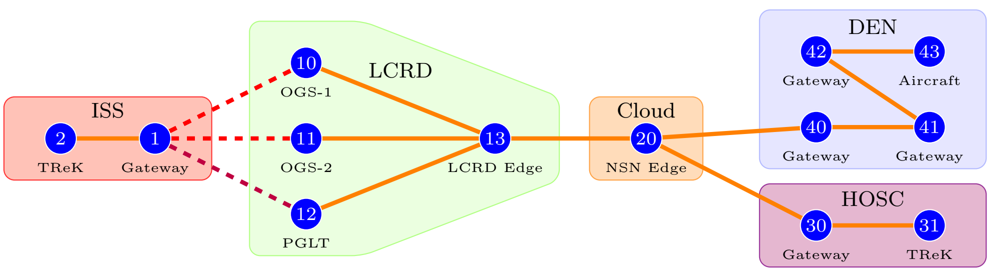

Figure 1: BP-layer network diagram: Omniscient

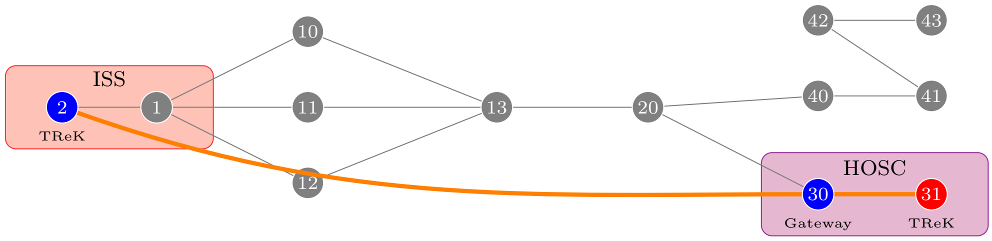

Figure 2: BP-layer network diagram: HOSC TReK Node

3. It inhibits debugging/troubleshooting.

Having one DTN node is also problematic:

1. It induces congestion.
2. It creates a single point of failure.
3. The ground infrastructure becomes a limiting factor.

Expanding on this last point, we note that while ILLUMA-T and LCRD could support a 1Gbps return link to either OGS, the Internet connection from the OGSs to the LCRD DTN nodes housed at White Sands Complex (WSC) provided 100-311Mbps service. In order to support testing at 1Gbps, a DTN node was installed at OGS-1 and DTN was used to overcome the link asymmetries[2].

Practically speaking, the ground station nodes are not destinations for bundle traffic: they function exclusively as routers. This enabled the true architectural contribution of these experiments: the contact plan for Node 1 viewed the LCRD network as a virtual node, Node 14; see Figure 3. From the perspective of Node 1, there are two adjacent bundle hops: Node 2 (the ISS TReK node) and Node 14. Any Earth-bound bundle from Node 1 must go through the LCRD network, but the next hop is not encoded in the bundle. Therefore, any such bundle will simply be routed to ILLUMA-T (and then to LCRD) as the IP layer, and whichever ground station is active will receive it. The other components of the ground network also routed traffic to Node 14, which improves scalability. Indeed, the TReK node at the HOSC (Node 31) also should not depend on knowledge of which ground station is being used when, and in fact with DTN it should not care if any link is established with the ISS at the time of transmission; see Figure 4 to see a depiction of the contact plan from the perspective of the cloud gateway (Node 20). LCRD operators would then bring up or down the appropriate ground station node (10, 11, or 12) depending on the LCRD schedule or handover, without levying any coordination requirements on the other users.

In order to realize this architecture, an LTP workaround had to be used. While LTP does not have addresses, it does have Engine IDs. For Node 1 to have only one LTP configuration it was necessary to give nodes 10, 11, and 12 the same Engine ID[3].

6 Tests and Results

In a very practical sense, there were two primary tests. One is that DTN must not impose a bottleneck on the system (which was physically limited to 1Gbps). Indeed, a networking solution that curtails the data yield is not a solution. It is known that HDTN and its LTP implementations can greatly surpass 1Gbps rates in software with a wide range of bundle sizes and RTTs[10]. The other primary test is that the architecture design with the virtual nodes shown in Figures 3 and 4 must work.

The DTN tests designed to prove out these higher-layer goals were as follows[4]:

- Bundle Protocol versions
  1. BPv6, with custody[5]
  2. BPv7, with and without BPSec
- Data
  1. Send small files (hundreds of 1-3MB files)
  2. Send large files (several 1-5GB files)

[2]See lessons-learned in Section 7.

[3]This leads to another lesson-learned; see Section 7.

[4]The tests and performance characteristics will be detailed in a dedicated upcoming paper, title TBD.

[5]See lessons-learned; Section 7.

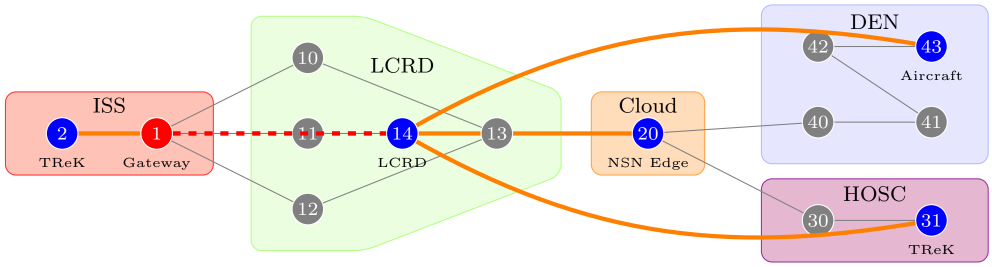

Figure 3: BP-layer network diagram: ISS Gateway

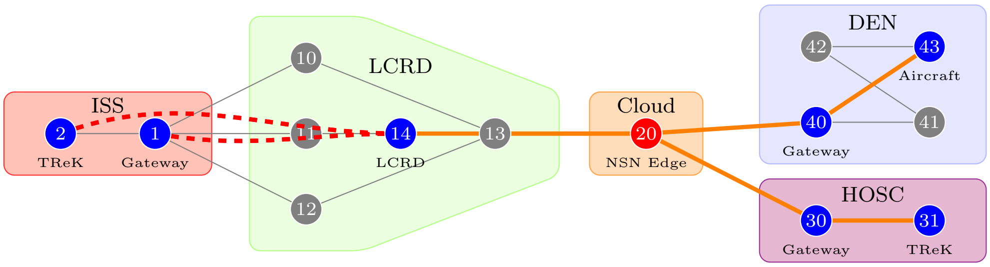

Figure 4: BP-layer network diagram: NSN Edge Node

3. Stream video

- Optical links
  1. Demonstrate that reliability at different layers improves data return over optical links
  2. Prevent data loss during handovers
- Network heterogeneity
  1. Space systems
  2. Ground systems
  3. Aerial systems
  4. Cloud systems
  5. Interoperability

6.1 BPv6

To demonstrate BPv6 with custody, small files (on the order of 2MB) were sent between Nodes 1 and 31. These tests involved three DTN implementations, HDTN v1.3, DTNME v1.1.0, and ION v3.6.1 (as part of TReK). No other BPv6 tests were conducted.

6.2 BPv7

The BPv7 tests used HDTN v1.3 exclusively, and all tests featured Node 1 as either a source or a destination with the other corresponding end either at Node 20 or Node 43 (for debugging purposes, Node 13 was occasionally used).

6.2.1 File transfer: Using the HDTN file transfer utilities, batch transfers of pictures (small files) and random data[6] (large files) were conducted repeatedly. All BPSec permutations (with or without integrity, with or without confidentiality) were run. Handovers were occasionally forced by LCRD operators during DTN transfers. Peak rates of 981Mbps were achieved.

6.2.2 Streaming: Video was streamed between Nodes 1 and 43, with each as source and destination[7]. 4k video files were streamed in both directions, and a live stream was conducted from Node 43 to Node 1.

Between configuration bugs and mismatches, pointing issues, schedule conflicts, etc., each test had to be run multiple times across the LCRD passes afforded this test campaign. Ultimately, every test was successfully conducted at least once and all goals were achieved.

7 Conclusion

The building blocks of the SSI must feature many dimensions of maturity, which includes the ability for these underlying protocols to support architectural requirements. In particular, it is necessary to subdivide the network into subnetworks that are largely isolated (self-contained). The test campaign conducted showed that DTN can be used to explore the concept of space network service provision.

In the process, several observations were made:

1. In some cases - depending on infrastructure and service requirements - DTN should be installed at ground stations.
2. While the ability to cut out intermediary nodes/subnetworks in a contact plan, configuration complexity would be greatly reduced by having a *default route*. Referring to Figure 2, Node 31 could have simply had a default route to Node 30. The notion of a default route is not free of subtlety, but warrants further research.
3. Giving multiple nodes the same LTP engine ID does not scale well. In our case the service provision was to one node (Node 1), but configuration management with multiple users would become complex. Instead, a formal approach to broad/multi/anycast at the BP layer would have provided a scalable solution.
4. In previous tests it was shown that the reliability of LTP enabled most ($>$ 95%) of bundles transmitted of the optical link from the aircraft to its ground station to be received, and custody transfer brought this number to 100%[10]. Custody transfer for BPv7 is both necessary and noticeably missing.'

7.1 Next Steps

With LunaNet in mind there are multiple obvious evolutionary steps:

- Create a network with multiple simultaneous users of the LCRD service provider. Our tests did feature

[6] pulled from `/dev/urandom`.

[7] For information on video streaming over DTN, see [17].

simultaneous transfers, but this was not thoroughly pushed.

- Extend the network to have multiple service providers.
- Combine the two above bullets.
- Research and develop the aforementioned scalability, configuration, and reliability shortcomings.

8 References

[1] V. Cerf, S. Burleigh, A. Hooke, *et al.*, "RFC 4838, Delay-Tolerant Networking Architecture," *IETF Network Working Group*, 2007. [Online]. Available: https://tools.ietf.org/html/rfc4838.

[2] S. Burleigh, K. Fall, and E. J. Birrane, *Bundle Protocol Version 7*, RFC 9171, Jan. 2022. DOI: 10.17487/RFC9171. [Online]. Available: https://www.rfc-editor.org/info/rfc9171.

[3] E. J. Birrane and K. McKeever, *Bundle Protocol Security (BPSec)*, RFC 9172, Jan. 2022. DOI: 10.17487/RFC9172. [Online]. Available: https://www.rfc-editor.org/info/rfc9172.

[4] D. J. Israel, K. D. Mauldin, C. J. Roberts, *et al.*, "Lunanet: A flexible and extensible lunar exploration communications and navigation infrastructure," in *2020 IEEE Aerospace Conference*, 2020, pp. 1–14. DOI: 10.1109/AERO47225.2020.9172509.

[5] M. Palsson, "Laser communications relay demonstration: Experiments with delay tolerant networking," English, *IET Conference Proceedings*, 106–111(5), Jan. 2023. [Online]. Available: https://digital-library.theiet.org/content/conferences/10.1049/icp.2024.0830.

[6] D. J. Israel, B. L. Edwards, J. D. Moores, S. Piazzolla, and S. Merritt, "The laser communications relay demonstration experiment program," *Ka and Broadband Communications Conference*, Oct. 2017.

[7] B. L. Edwards and D. J. Israel, "Update on nasa's laser communications relay demonstration project," in *2018 SpaceOps Conference*. DOI: 10.2514/6.2018-2395. eprint: https://arc.aiaa.org/doi/pdf/10.2514/6.2018-2395. [Online]. Available: https://arc.aiaa.org/doi/abs/10.2514/6.2018-2395.

[8] A. Seas, B. Robinson, T. Shih, F. Khatri, and M. Brumfield, "Optical communications systems for NASA's human space flight missions," in *International Conference on Space Optics — ICSO 2018*, Z. Sodnik, N. Karafolas, and B. Cugny, Eds., International Society for Optics and Photonics, vol. 11180, SPIE, 2019, pp. 182–191. DOI: 10.1117/12.2535936. [Online]. Available: https://doi.org/10.1117/12.2535936.

[9] T. Basciano, B. Pohlchuck, and N. Shamburger, "Application of delay tolerant networking on the international space station," *ISS Research and Development Conference*, Jul. 2019.

[10] A. Hylton, J. Cleveland, R. Dudukovich, *et al.*, "New horizons for a practical and performance-optimized solar system internet," in *2022 IEEE Aerospace Conference*, 2022.

[11] NASA, *High-rate delay tolerant networking (hdtn)*. [Online]. Available: https://www1.grc.nasa.gov/space/scan/acs/tech-studies/dtn/.

[12] M. Ramadas, S. Burleigh, and S. Farrell, "RFC 5326, Licklider Transmission Protocol - Specification," *IETF Network Working Group*, 2008. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc5326.

[13] K. Scott and S. Burleigh, "RFC 5050, Bundle Protocol Specification," *IETF Network Working Group*, 2007. [Online]. Available: https://tools.ietf.org/html/rfc5050.

[14] A. Hylton, B. Mallery, J. Hwang, *et al.*, "Multi-domain routing in delay tolerant networks," in *2024 IEEE Aerospace Conference*, 2024, pp. 1–20. DOI: 10.1109/AERO58975.2024.10521176.

[15] J. A. Fraire, O. De Jonckère, and S. C. Burleigh, "Routing in the space internet: A contact graph routing tutorial," *Journal of Network and Computer Applications*, vol. 174, p. 102 884, 2021, ISSN: 1084-8045. DOI: https://doi.org/10.1016/j.jnca.2020.102884. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1084804520303489.

[16] S. Booth, R. Dudukovich, N. Kortas, *et al.*, "High-rate delay tolerant networking (hdtn) user guide," NASA, Glenn Research Center, Cleveland OH 44135, USA, Technical Memorandum NASA/TM-20230000826/REV1, May 2024.

[17] K. J. Vernyi and D. Raible, "4k high definition video and audio streaming across high-rate delay tolerant space networks," in *AIAA SCITECH 2024 Forum*. DOI: 10.2514/6.2024-1738. eprint: https://arc.aiaa.org/doi/pdf/10.2514/6.2024-1738. [Online]. Available: https://arc.aiaa.org/doi/abs/10.2514/6.2024-1738.