LASER COMMUNICATIONS RELAY DEMONSTRATION: EXPERIMENTS WITH DELAY TOLERANT NETWORKING

Alan Hylton[1], David Israel[1], Mikael Palsson[2]
alan.g.hylton@nasa.gov

Keywords: DTN, LCRD

Abstract

The Laser Communications Relay Demonstration (LCRD) is a NASA mission that is presently providing a link between two optical ground stations on Earth and geosynchronous orbit; this capability will be extended to include a laser terminal on the International Space Station (ISS). From a practical perspective, LCRD provides a 1Gbps Ethernet link with an anticipated round-trip time of 4 seconds. Delay Tolerant Networking (DTN) is being used to transcend LCRDs capabilities from a point-to-point link to true network connectivity. Indeed, integrating the LCRD link into any network, such as one aboard the ISS, adds new hops and paths. In this paper, we detail DTN experiments conducted using LCRD to demonstrate how DTN and optical capabilities complement each other to enable the next generation of space communication modalities. We conclude with considerations on using LCRD to connect the ISS to the optical ground stations.

1 Introduction

The Laser Communications Relay Demonstration (LCRD) represents NASA's intention to advance optical communications technology and its infusion into both deep-space and near-Earth operations. Intrinsically, LCRD provides bidirectional optical communications between geosynchronous Earth orbit (GEO) and Earth [1]; we will provide a brief overview. Extrinsically, LCRD adds a link to existing and future systems. If there are extant communications channels, it adds new hops and new paths to existing infrastructure. The purpose of this paper is to detail networking experiments conducted over LCRD to date, lay out a roadmap of upcoming tests, and to discuss future directions.

1.1 LCRD

LCRD is hosted by the Space Test Program Satellite 6 (STPSat-6), which launched on December 7th of 2021 [2]. The payload includes two optical communications terminals, which can operate at 1.244Gbps, and the spacecraft augments the capability with a Ka-band system (Transmit: 622 Mbps, Receive: 64Mbps). LCRD is supported by two optical ground station (OGS), one at Table Mountain in California (OGS-1), and the other on Haleakala in Hawaii (OGS-2). RF support is provided at the White Sands Complex. As LCRD has two optical terminals, it can function as a relay between OGS-1 and OGS-2, or it could be a relay between an asset in low Earth orbit (LEO) and one OGS. In the latter, LCRD might choose between an OGS given weather considerations, e.g. cloud cover [3].

As a prelude to deeper networking considerations, note that the OGS site prioritization for cloud-free line of sight (CFLOS) comes at the expense of terrestrial infrastructure, which must be made symmetric to the LCRD capabilities. Indeed, the utility of a 1Gbps connection to geostationary orbit (GEO) or low earth orbit (LEO) is diminished if the data are lost on the ground. This does not mean, however, that the connection (Internet or otherwise) from the OGSs must be 1Gbps or higher. Rather, a store, carry, and forward approach could be taken to buffer data on the ground for rate-matching. One goal of networking would be to not only provide such a mechanism, but to provide such a mechanism that is standardized.

1.2 Networking

The Union of Concerned Scientists estimate that as of January 1st, 2022, there were over 4,852 operational satellites about the Earth [4]. In May of 2022, it was estimated that of the then total, 2,216 were operational Starlink satellites[5]. These figures are unprecedented, and as satellites of all types continue to grow in number so to must the overall communications capability. The particular capability needed is *scalability*, which means developing, implementing, and operating a meaningful, standardized networking technology. The nature of space communications makes this difficult, for example, due to:

1. the lack of end-to-end connectivity,
2. the high latencies and their variances, and
3. the overall system is highly mobile.

The approach taken by NASA is Delay Tolerant Networking (DTN)[6], an overlay network protocol that is designed to enable communications in these disconnected networks. While research continues into the underlying nature of DTNs (e.g. [7]), several DTN implementations have matured sufficiently for operational consideration, which can be informed via LCRD experiments. This includes the Interplanetary Overlay Network (ION), High-rate DTN

(HDTN), and DTN Marshall Edition (DTNME).

A future goal of this current research is to understand how to use LCRD to enable more expansive networking aboard the International Space Station (ISS). Indeed, the first low Earth orbit (LEO) user of LCRD will be Integrated Laser Communications Relay Demonstration Low-Earth Orbit User Modem and Amplifier Terminal (ILLUMA-T), which will be integrated into the International Space Station (ISS) [8]. ILLUMA-T, combined with LCRD, will create an additional path to the gigabit Ethernet network aboard the ISS and the ground stations, which extends to its extant, operational ION-based DTN [9]. Keeping this future goal in mind, we will outline DTN, tests (both conducted and scheduled), and the future ILLUMA-T campaign.

2 DTN

The networking standard, DTN, is designed to address the aforementioned generalized nature of networking in space. In a standard terrestrial network, typical actions upon receiving a packet are forwarding and dropping. By extending the actions to store, carry, and forward data, DTN enables routing approaches that operate regardless of end-to-end connectivity at any instant. Typically, this is done with schedule-based routing; for the tests conducted here, the tests are run when LCRD operations allows it, and hence the schedule is simply "up."

In a DTN, the primary unit of data is a *bundle*. For our considerations, the important features are that a bundle can be of essentially any size, and that there is optional reliability. To transfer a bundle from one node to another, a user could use familiar protocols such as the Internet Protocol (IP), or namely the Transmission Control Protocol (TCP) or User Datagram Protocol (UDP). We see then that DTN is an *overlay* network, providing a high-level network where typical network functions (e.g. routing, management, and security) are managed by DTN, and what appears to be the atomic point-to-point connections between DTN nodes are really underlying networks. As described, the benefit is limited: it is fair to ask what place TCP has here. The answer is that for some hops, such as those between ground stations, IP makes sense. Moreover, using the `send` system call, an arbitrarily-sized bundle can be sent easily. Looking a little deeper, we realize that TCP routing assumes end-to-end connectivity, and DTN segments this according to local availability which is determined by schedules. However, this is taken a step further by implementing additional protocols built around whichever particular requirements are relevant. This includes the Licklider Transmission Protocol (LTP) [10], which is a reliable protocol with tunable timers for reliability of links with *any* one-way light time (OWLT). DTN can be extended to send bundles over any protocol, and a shim layer known as a *convergence layer* is used to translate the bundles and the underlying layers accordingly (in particular most frames cannot be arbitrarily large). The takeaway is that DTN turns disparate collections of links into a system, and hence is the prime candidate for achieving a positive returns to scale with communicating assets in space.

Although the distance to GEO is roughly 125ms (for a 250ms round trip time (RTT)), which could be force-fit with TCP, there are additional delays to consider. A particular example comes from the interleavers used to overcome intermittant loss. Overall, LCRD is designed to operate with a 4s RTT, precluding typical Internet protocols. For high latency lunar and deep space links, LTP is the protocol of choice for reliable transport over CCSDS space link protocols. By combining transport reliability with DTN's network-layer reliability, using *custody transfer*, DTN can optimizes flows over the asymmetrical links in the overall system.

The standards for DTN have given rise to numerous software implementations, each with their own design criteria and characteristics. Below we briefly recall the features of the implementations of interest.

2.1 ION

ION has been NASA's reference implementation for DTN, developed by JPL. ION was used for the first LCRD DTN tests to demonstrate the functionality of DTN over the optical link. ION was designed for operations in deep-space, and hence is not optimized for performance but rather for adherence to a particular software standard for embedded development.

2.2 HDTN

High-rate DTN (HDTN)[11] is a performance-optimized implementation of DTN, has been demonstrated over analogous laser links, albeit with aircraft[12]. This test included the transmission of 34GB worth of image data over an intermittent, free space optical link, and was able to saturate the roughly 900Mbps link with LTP [13]. HDTN will be used on the ISS as a high-speed DTN gateway, and is designed to run on bare metal or in containers.

2.3 DTNME

DTN Marshall Enterprise (DTNME)[14] is the third implementation of DTN considered, written as a modernization and overhaul of the older DTN2 software. DTNME is currently operating in the Huntsville Operations Support Center (HOSC) and the ISS.

2.4 DTN

As alluded to, the ISS currently features an operational DTN. This includes ION and DTNME nodes, and is being extended to include HDTN. Because these software packages are all implemented from the same specification, they can interoperate. This means that low-powered embedded data-sources (leaf nodes) could use ION, connected to an on-board gateway running HDTN to opportunistically forward bundles at line-rate to the ground, where bundles are then forwarded to their respective project PIs. Hence all

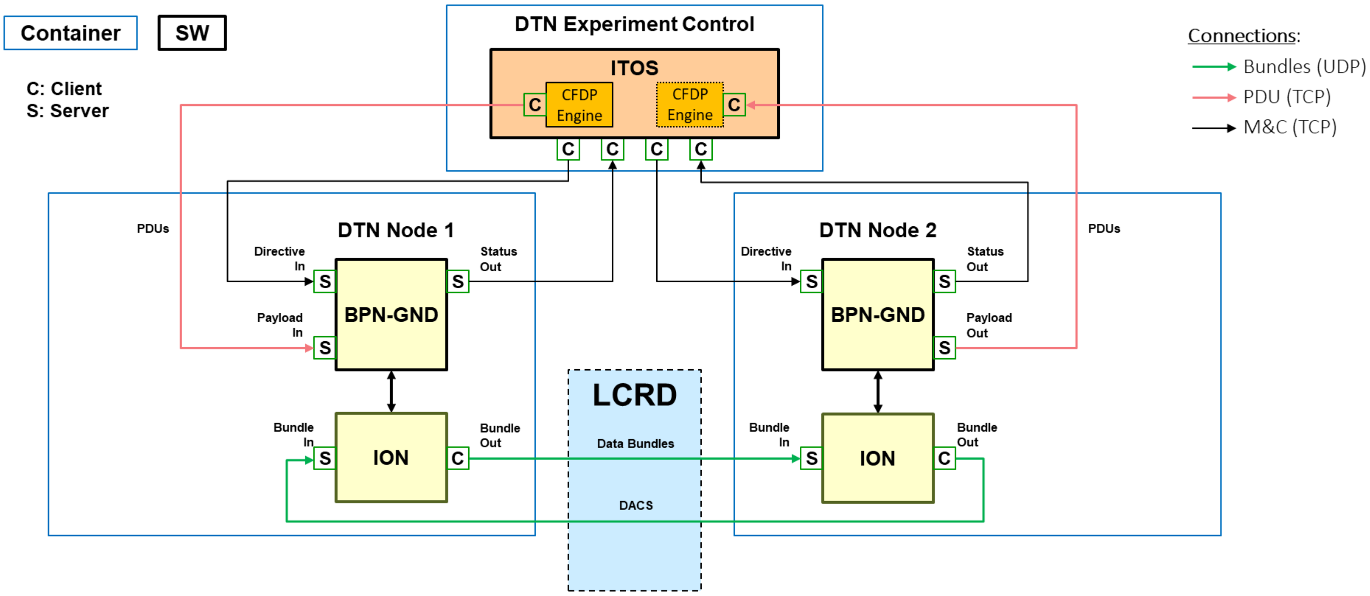

Figure 1: LCRD experiment from DTN perspective

the data formats are compatible, and a measure of heterogeneity is allowed. Interoperability also implies performance metrics must also make sense. At the lower layers we might consider bits or symbols per second, building up to packets per second. The bundle data type comes with non-trivial processing times, and bundles per second must also be considered. Indeed, the different DTN implementations are separated by several orders of magnitude in terms of bundles per second.

3 ION Tests

3.1 *Environment*

LCRD can be configured for many different link types, and in this case, it relayed data from OGS-1 back to OGS-1. It should be noted that from an Ethernet perspective, this effectively reduces LCRD to a media converter with a 4s RTT. The first tests using containerized ION nodes had a data as WSC ↔ OGS-1 ↔ LCRD ↔ OGS-1 ↔ WSC. The controller used was the Integrated Test and Operations System (ITOS)[15], which is a telemetry and command system that can be used for testing, development, simulation, and operations. In our case, it interfaced with ION via a glue-layer known as BPN-GND (the details of which are not necessary for this discussion; see [16] for details). The data were files, which were transmitted using non-reliable CCSDS File Delivery Protocol (CFPD). Note that reliable CFDP was not used as reliability was ensured by DTN custody transfer; as a matter of notation, custody signals are often aggregated for performance reasons, and these aggregates are called DTN Aggregate Custody Signals (DACS). This is illustrated in Figure 1.

The optical ground stations are connected via Ethernet networks to the White Sands Complex (WSC), where the optical data processors reside in the so-called Experimenter's Rack. The security posture is such that the tests are operated from the NASA Goddard Space Flight Center (GSFC), but conducted at WSC. The experimenter's rack contained a PC that handled the raw LCRD data - software colloquially known as HEIDIs run in virtual machines - and Docker. ION was configured to transmit bundles over UDP as opposed to LTP to ease configuration. This perspective is shown in Figure 2.

3.2 *Configurations and environment*

The test operating characteristics were chosen to demonstrate "weather fade". Fifty files at 16.8MB each were sent to demonstrate DTN's robustness against optical link degradation. Each test was originally expected to last roughly 30 minutes; this includes the payload, ground station, and HEIDI setup times as well as closeout times.

It should be noted that CFDP can be configured to segment files into chunks with a maximum size, and these chunks can be put in corresponding bundles that fit within a UDP datagram. In this case, the maximum protocol data unit (PDU) size was limited to 48,000 bytes, the maximum allowed by ITOS.

DACS transmission rates are also configurable, and can have an impact on performance; to minimize this, the lowest allowed setting of 1-second transmission rates was chosen. The retransmission time was set to 10 seconds, which is driven by the LCRD RTT and the DACS transmission rate.

The test was run three times, and although the weather was reported as "fair" for all three runs, the tests were operated in the daytime and local noon is the worst time of day from a noise perspective.

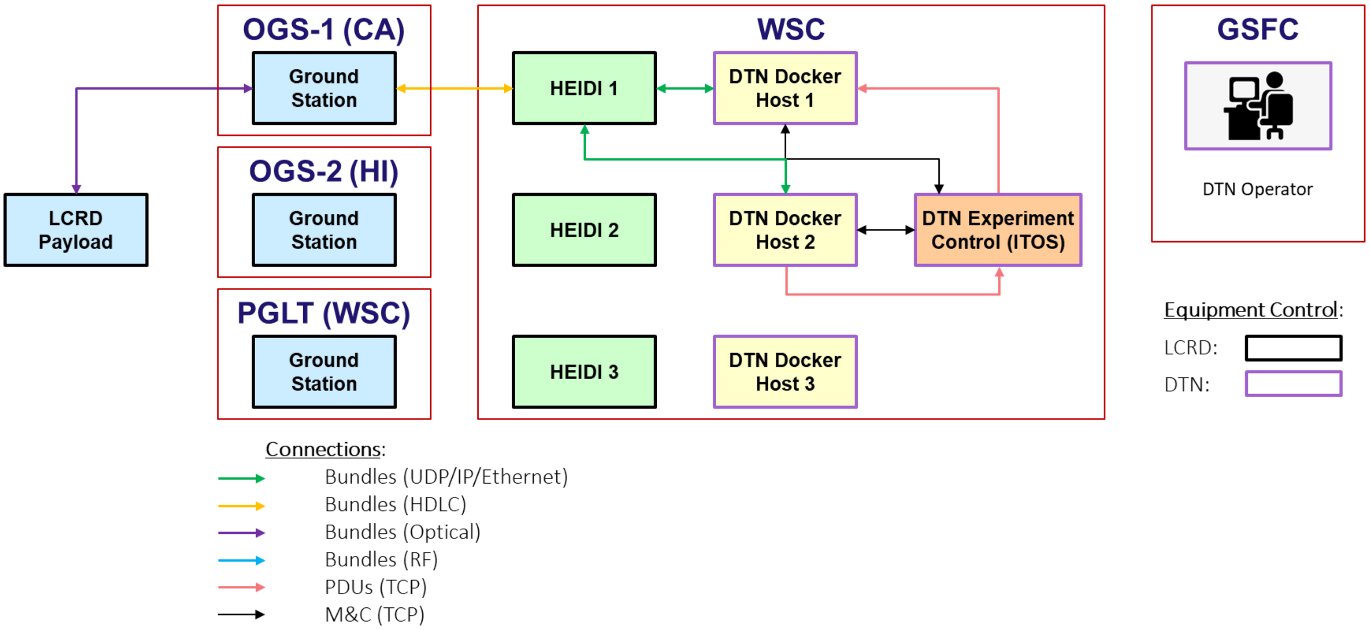

Figure 2: LCRD experiment environment

3.3 *Observations*

The highlight of the tests is that all data made it from their source to their destination - eventually. However, significantly more bundles were lost than successfully received. At a glance, the results are shown in Table 1. Here the bundles sent are data bundles as opposed to DACS.

The tests run were in no way expected to push the performance of LCRD. Indeed, neither is ION optimized for performance as a software implementation nor was it configured for performance. Moreover, the rack resources were limited, as the DTN nodes, the HEIDIs, and various security and auditing software were all running on the same physical hardware. Rather the purpose was to show that in the face of a large and variable bit error rate (BER), DTN was able to provide robustness to data transfer. In this sense, ION was highly successful. It was interesting to note that the uplink BER was generally worse than the downlink BER, though as expected there was a close correlation between code error rate and data flowing through the link (High-level Data Link Control (HDLC) counters versus bundle counters).

For a time-series representation, Runs 1 and 2 of Table 1 are graphed in Figure 3. On the left hand side is the number of bundles buffered waiting for custody acceptance - this grows during loss when the DACS are not successfully transmitted. The unstable nature of the link can be see visually for the first run. The corresponding bit rate from the file transfer application viewpoint is shown on the right hand side. It is apparent that the link flapping was less pronounced in Run 2, which corresponds with the data in the table.

4 Future tests

4.1 *Upgrades*

Following the successful ION tests, the hardware in the Experimenter's Rack was upgraded to modern NUCs and Ethernet switches. The purpose of the new hardware is to allow the HEIDIs to operate while sparing enough cycles to also run the DTN network stack for everything to run at line-rate. The new hardware has been integrated into the rack, and following integration tests will support the next round of testing.

4.2 *HDTN and DTNME*

Having demonstrated the utility of DTN, the next phase is to show that DTN does not impose a bottleneck in the system. This is considered crucial for widespread DTN adoption. To this end, a collection a high-rate tests are scheduled. The HDTN and DTNME nodes will also be deployed in Docker containers. Each test will focus on file transfer over a pair of these implementations. For this generation of tests, the nodes will send bundles from one to the other using LTP/UDP/LCRD, marking another departure from the first tests. The data source will be roughly 34GB of files, which will be transferred and check summed albeit not with CFDP but rather a performance optimized file transfer utility written specifically for DTN.

The test parameters also include transfer with and without custody. For the first tests, custody was necessary as UDP is unreliable and CFDP was specifically configured to run unreliably as well. Note that in [12] for 100% success it was necessary to couple reliability at various layers.

The objectives are to

- gain greater insight into system performance, limitations, and functionality;

Run	Fwd HDLC Frames	Rtn Good Frames	Lost Frames	Bundles Sent	Bundles Rcvd	Resent Bundles
1	5,178,480	1,364,801	3,813,679	157,558	20,953	136,605
2	1,250,903	923,406	327,497	102,088	20,135	81,953
3	3,142,272	1,690,838	1,451,434	95,561	23,582	71,979

Table 1: Statistics of the three test runs

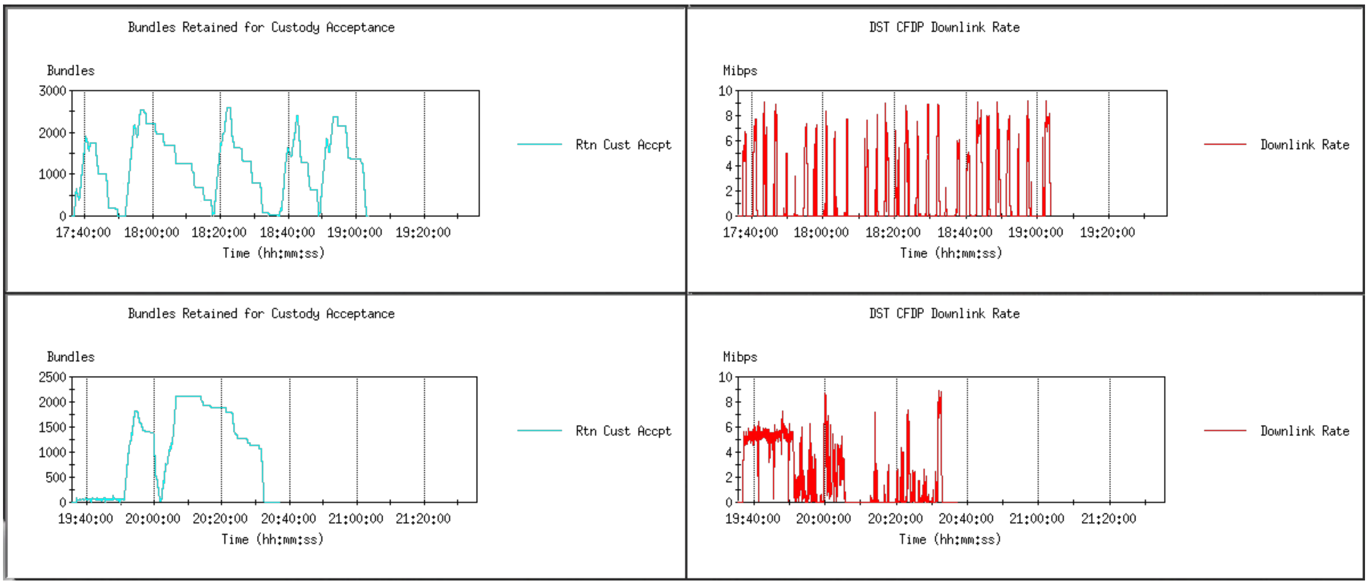

Figure 3: LCRD test results

- provide insight to inform architecture and implementation trades, including the DTN stack and Network Management, as well as ground system integration strategies; and
- showcase the potential for the LCRD platform to be utilized not just for optical and link experiments, but also for network and routing style experiments.

This last item highlights that LCRD's ability to extend DTN and DTN operations towards NASA's upcoming LunaNet[17]; before LunaNet, the ILLUMA-T tests must be conducted, and prior to that a third round of LCRD tests will be conducted.

4.3 Pre-ILLUMA-T

As mentioned, there is a DTN operational aboard the ISS. The upcoming ILLUMA-T payload will allow the ISS to establish an additional link to its existing RF capabilities to the ground via LCRD. Before ILLUMA-T and the ISS can be integrated into the ground network, the capabilities of the ground network must be extended to include both greater network connectivity and network services.

For connectivity, we must consider how data received at an OGS (presumably at 1Gbps) would be forwarded to their intended recipient(s). Both the links between the OGSs and WSC must be considered; returning to the rate-mismatch scenario from the introduction, DTN (and hence buffering) would be required at WSC. Moreover, the security posture of WSC could make it difficult for partners at NASA field centers to connect directly. Both of these concerns can be addressed by flowing high-rate bundle data from a DTN node at WSC through LCRD to a receiver in a secure cloud platform. Indeed, it is a goal of the LCRD Guest Experimenters's Program to make LCRD testing accessible to those outside of the WSC boundary[18]. Hence it is reasonable to extend the DTN network from LCRD to such a cloud platform, which would enable future missions (such as ILLUMA-T) to further extend the network to the ISS.

The test plan is then to begin by transmitting files in bundles from the Experimenter's Rack to the cloud. This test can be modified in several ways, such as to include security (BPsec [19]) or to replace the files with video streaming, a new feature found in HDTN, which will allow more performance metrics than bit/bundle rates, such as various streaming quality metrics.

These tests are expected to be run prior to the launch of ILLUMA-T.

5 Conclusion

LCRD as an operational platform allows for the first-of-their-kind DTN tests using optical relay satellites. As DTN aggregates such links into a greater network, large high-performance networks can start to be studied to determine feasible paths for large DTN deployment and integration.

The first tests used ION, and despite older hardware and link flapping successfully demonstrated reliability due to DTN in file-transfer applications. This marks the first of several DTN tests. Now that the hardware at WSC has been upgraded, the next tests will feature HDTN and DTNME to complete a trade on more performant options for infusion consideration. Following the first wave of high-rate DTN tests, a more true networked scenario featuring multiple DTN hops and cloud integration will take place, offering a path from LCRD to NASA field centers. The goal is to create a ground-based infrastructure that NASA can use to ultimately connect to the ISS via ILLUMA-T.

Thinking beyond ILLUMA-T, NASA's plans for the Lunar network, LunaNet [17], include a multi-hop multi-path network that supports DTN. By incrementing towards this level of complexity, any gaps in DTN can be identified and closed as experienced is gained with its operations.

6 References

[1] 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.

[2] L. Mohon, *Laser communications relay demonstration (lcrd)*, Jan. 2022. [Online]. Available: `https://www.nasa.gov/mission_pages/tdm/lcrd/index.html`.

[3] 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`.

[4] U. of Concerned Scientists, *Ucs satellite database*, Jan. 2022. [Online]. Available: `https://www.ucsusa.org/resources/satellite-database`.

[5] J. McDowell, *Starlink statistics*, Oct. 2021. [Online]. Available: `https://planet4589.org/space/stats/star/starstats.html`.

[6] 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`.

[7] R. Short, A. Hylton, J. Cleveland, *et al.*, "Sheaf theoretic models for routing in delay tolerant networks," in *2022 IEEE Aerospace Conference*, 2022.

[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] 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`.

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

[12] 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. [Online]. Available: `https://ntrs.nasa.gov/citations/20220003634`.

[13] ——, "New horizons for a practical and performance-optimized solar system internet," in *2022 IEEE Aerospace Conference*, 2022.

[14] NASA, *Dtn marshall enterprise (dtnme)*. [Online]. Available: `https://software.nasa.gov/software/MFS-33978-2`.

[15] ——, *Integrated test and operations system (itos)*. [Online]. Available: `https://itos.gsfc.nasa.gov/`.

[16] ——, *What are some dtn software options?* [Online]. Available: `https://www.nasa.gov/directorates/heo/scan/engineering/technology/disruption_tolerant_networking_software_options`.

[17] 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`.

[18] D. J. Israel, B. L. Edwards, R. L. Butler, *et al.*, "Early results from NASA's laser communications relay demonstration (LCRD) experiment program," in *Free-Space Laser Communications XXXV*, H. Hemmati and B. S. Robinson, Eds., International Society for Optics and Photonics, vol. 12413, SPIE, 2023, p. 1 241 303. DOI: `10.1117/12.2655481`. [Online]. Available: `https://doi.org/10.1117/12.2655481`.

[19] 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`.