

Spacecraft Data and Relay Management Using Delay Tolerant Networking

Christopher J. Krupiarz*

The Johns Hopkins University Applied Physics Laboratory, Laurel, MD, 20723-6099

Esther H. Jennings[†], Jackson N. Pang[‡], Joshua B. Schoolcraft[§], John S. Seguí^{**}, and J. Leigh Torgerson^{††}
Jet Propulsion Laboratory, Pasadena, CA, 91109

NASA's demonstration of the successful transmission of relay data through the orbiting Mars Odyssey, Mars Global Surveyor, and Mars Express by the Mars Exploration Rovers has shown not only the benefit of using a relay satellite for multiple landed assets in a deep space environment but also the benefit of international standards for such an architecture. As NASA begins the quest defined in the Vision for Exploration with robotic and manned missions to the Moon, continues its study of Mars, and is joined in these endeavors by countries world-wide, landed assets transmitting data through relay satellites will be crucial for completing mission objectives. However, this method of data delivery will result in increased complexity in routing and prioritization of data transmission as the number of missions increases. Also, there is currently no standard method among organizations conducting such missions to return these data sets to Earth given a complex environment. One possibility for establishing such a standard is for mission designers to deploy protocols which fall under the umbrella of Delay Tolerant Networking (DTN). These developing standards include the Bundle Protocol (BP) which provides a standard, secure, store and forward mechanism designed for high latency and asymmetric communication links and the Licklider Transmission Protocol (LTP) which is used to provide a reliable deep space link transmission service.

As part of the Mars Technology Program being managed by the NASA Jet Propulsion Laboratory, a joint team of researchers at the NASA Jet Propulsion Laboratory (NASA/JPL) and the Johns Hopkins University Applied Physics Laboratory (JHU/APL) are studying DTN-developed protocols for use in a next generation Mars protocol architecture. Particularly, the team is developing flight software versions of BP and LTP defined in DTN as well as creating various test scenarios to be modeled and simulated through a NASA/JPL developed network test suite – Multimission Advanced Communications Hybrid Environment for Test and Evaluation (MACHETE). The team is simulating a mission scenario involving two landers, two relay orbiters, ground stations and a ground system using MACHETE to verify the feasibility of using the DTN in future missions.

Nomenclature

<i>AOS</i>	=	Advanced Orbiting Systems
<i>BP</i>	=	Bundle Protocol
<i>CCSDS</i>	=	Consultative Committee for Space Data Systems
<i>CFDP</i>	=	CCSDS File Delivery Protocol
<i>DSN</i>	=	Deep Space Network
<i>DTN</i>	=	Delay Tolerant Networking
<i>IP</i>	=	Internet Protocol
<i>IPN</i>	=	Interplanetary Internet
<i>LTP</i>	=	Licklider Transmission Protocol
<i>MACHETE</i>	=	Multimission Advanced Communications Hybrid Environment for Test and Evaluation
<i>MER</i>	=	Mars Exploration Rover

* Senior Professional Staff, Space Department, 11100 Johns Hopkins Road, Laurel, MD 20723

† Technical Staff, Communications Networks Group, 4800 Oak Grove Drive M/S 238-343, Pasadena, CA 91109

‡ Engineer, Communications Networks Group, 4800 Oak Grove Drive M/S 238-420, Pasadena, CA 91109

§ Engineer, Communications Networks Group, 4800 Oak Grove Drive M/S 238-420, Pasadena, CA 91109

** Engineer, Communications Networks Group, 4800 Oak Grove Drive M/S 238-343, Pasadena, CA 91109

†† Senior Engineer, Communications Networks Group, 4800 Oak Grove Drive M/S 238-420, Pasadena, Ca 91109

<i>MOC</i>	= Mission Operations Center
<i>PDU</i>	= Protocol Data Unit
<i>PTL</i>	= Protocol Test Lab
<i>SFO</i>	= Store-and-Forward Overlay
<i>SCPS</i>	= Space Communications Protocol Specification
<i>SLE</i>	= Space Link Extension
<i>SOAP</i>	= Satellite Orbital Analysis Program
<i>TCP</i>	= Transmission Control Protocol
<i>TC</i>	= Telecommand
<i>TM</i>	= Telemetry
<i>UDP</i>	= User Datagram Protocol

I. Introduction

CURRENTLY at Mars there are the orbiters Mars Odyssey, Mars Global Surveyor, Mars Express, and Mars Reconnaissance Orbiter and the two Mars Exploration Rovers, Spirit and Opportunity. The rover vehicles return up to 95% of their science data via relays through these orbiters as opposed to a Direct-to-Earth link. Mars Odyssey is used as the primary relay for this purpose with Mars Express having been used as a demonstration of an international relay and Mars Global Surveyor being available for emergencies. The relay spacecraft have a circular buffer in which data packets from the rovers are stored. This data is stored on the same Solid State Recorder as the science data for the orbiter or, as in the case of Mars Global Surveyor, within one of the instrument's storage. The data is then downloaded when bandwidth is available on the Deep Space Network.

While the extensive amount of data returned from the rovers has demonstrated this architecture to be highly successful and functional, as the number of assets both on Mars and in orbit increases, the complexity of preplanning all the communication paths of these missions has the potential to be overwhelming. In fact, the planning process may be described as a constrained optimization problem where communication links may be unavailable for long periods of time and a finite storage constraint exists at each asset¹; the optimization objective might be to minimize the average delay of data from sender to receiver.

Furthermore, as organizations outside of the NASA Jet Propulsion Laboratory (NASA/JPL) teams operate spacecraft at Mars, mission planners will find it problematic to plan across disparate organizations. From a system perspective, this architecture precludes the ability of the landers to delete data until it is confirmed on the ground (which can take many minutes or hours due to round trip light time, line of sight, or operational delays), does not allow for prioritizing of data across missions and instruments, and does not provide a well-documented standard by which future missions can "plug-and-play" into the Mars network infrastructure.

In response to a Research Announcement in 2004, NASA awarded a contract to The Johns Hopkins Applied Physics Laboratory (JHU/APL) and NASA/JPL to develop a suite of "next generation" protocols for future Mars missions. The goal of this project is two-fold: (1) to model and simulate future Mars networks and potential protocols and (2) to develop flight software implementations of selected protocols. Along with the current suite of protocols defined by the Consultative Committee for Space Data Systems (CCSDS), the team selected additional developing standards, including the Licklider Transmission Protocol (LTP) and the Bundle Protocol (BP), that fall under the Delay Tolerant Networking (DTN) umbrella. The responsibilities on this project are divided between NASA/JPL performing the simulations with input from JHU/APL and JHU/APL engineers developing the flight software.

II. Delay Tolerant Networking

DTN began as an effort to standardized communications for the Interplanetary Internet (IPN)². As work progressed, researchers observed that deep space communications was, in fact, a specialized domain of a larger group of challenged networks. These networks all had similarities in that they experienced several of these features: asymmetric communication, high error rate links, long delays, and intermittent connectivity. As a result, the network community is developing a body of research and both NASA and the Defense Advanced Research Projects Agency (DARPA) have established funding to continue work on this technology.

LTP is a long haul protocol which provides a reliable service across a deep space link³. Evolved from the retransmission procedures of the CCSDS File Delivery Protocol (CFDP)⁴, LTP ensures that data sent from a spacecraft to the ground or vice versa is received correctly without the additional overhead of manual verification of the data. BP is a standard protocol for performing store-and-forward transmission and is an overlay network which provides a common interface handling data units called bundles across heterogeneous networks while maintaining an end-to-end networking capability⁵. BP acts as a common overlay for hosts running disparate transport layers which allow tuned protocols to be used in regions where they are appropriate. Through this BP overlay, these different transport protocols that couldn't communicate with each other before will be able to do so. This is essential for a Mars network as typically data from a rover or lander is not sent through a real-time relay link. Instead, data must be stored on the orbiter to await transmission and it must rely on different underlying protocols at different points along the communication path. In a Mars relay environment, BP and LTP would be used in conjunction with each other. A sample protocol stack for BP/LTP is shown in Fig. 1. This example shows how LTP and BP can be inserted into a standard protocol stack with BP as the common protocol across networks using three different underlying protocols: 1) the lander reliably transmits bundles via the CCSDS Proximity-1 protocol to the orbiter, 2) bundles flow over LTP to reach Earth with guaranteed delivery, and 3) the bundles are transmitted over the Internet using standard terrestrial communication such as TCP/IP.

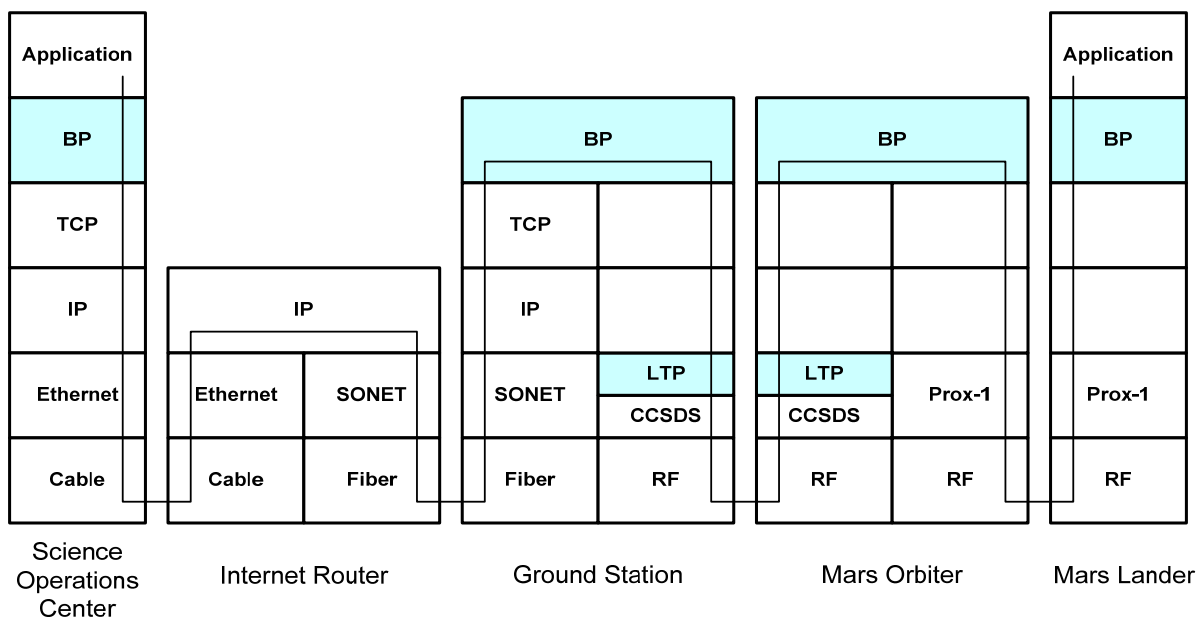


Figure 1. BP/LTP Protocol Stack for Future Mars Architectures

The Bundle Protocol operates at a sub-application “bundle” layer providing end-to-end communication over performance-challenged networks while allowing interoperability between highly heterogeneous networks such as IP does; however, through the use of specialized convergence layers adapts the bundle protocol “sits” on top of each local internet’s preferred transport protocol and is more flexible than IP allowing for interconnecting greater differing network types than IP. Analysis and simulation has shown that TCP does not scale to the requirements of the InterPlanetary Network due to the long round trip light times and high error rates. It is worth noting that BP does not replace IP but rather uses a TCP/IP convergence layer adapter when connection is needed on an IP-based network. For example BP would use TCP, LTP, CCSDS Proximity-1 Space Link Protocol convergence layer adapters for the terrestrial Internet, the interplanetary backbone, and the planetary surface networks, respectively.

III. Advantages of DTN Protocols

As part of its initial analysis of whether introducing BP/LTP in a Mars environment was worth the added complexity in flight software, the team examined the current MER Mars relay baseline architecture in an effort to determine what benefits BP/LTP could provide.

A. Standardization

Once the LTP and BP specifications are finalized, BP/LTP will provide a standard method of store-and-forward communication in deep space as well as other terrestrial applications. Researchers in the worldwide networking community are currently working towards standardization of both LTP and BP. Unlike proprietary or closed protocols, the development process for LTP and BP is open to contributions from experts not only in the space environment, but those from other areas of specialization. This process allows BP/LTP to be rigorously examined and evaluated by a large number of individuals as opposed to only a small number in a typical proprietary development environment. Finally, international standards will not limit communication between spacecraft of different organizations or countries.

B. Resource Savings

With current Mars operations, mission controllers only delete data from a landed asset upon confirmation that the data was received correctly on Earth. As a result, given round trip light time ranges at Mars of 10-40 minutes, a significant delay occurs from the time the data is sent from the lander to Earth and the time that the data is deleted. During this time, the data occupies valuable storage on-board the orbiter and landed assets. BP/LTP removes this requirement by providing reliable store-and-forward through custody transfer⁶. Once the data has been sent from the spacecraft, the protocols ensure reliable delivery to the mission and science operations centers.

C. Dynamic Data Prioritization

Currently, data are stored on-board a relay satellite with the current Mars network infrastructure via first-in, first-out queue with relay data taking precedence over orbiter spacecraft data. No dynamic prioritization is done regarding the actual contents of the data; prioritization is dependent on the scheduled command sequences. This shortcoming does not inherently permit expedited or emergency data from taking a higher priority when bandwidth is available to Earth. BP directly provides this capability through its prioritization scheme, which contains low/medium/high priority levels. Through this system, high value data can take precedence on an as needed basis.

D. Delay Tolerance

BP is designed to be delay tolerant. Bundles can be produced without concern for when they will be transmitted during a subsequent Deep Space Network (DSN) link opportunity. However, this may not always be sufficient in the event of time dependent data. BP accounts for this through the use of a time to live field. In the event that certain science or spacecraft data may become “stale” if the relay satellite does not deliver the data in a timely fashion, the data will be recognized as a candidate for deletion when storage becomes scarce.

E. Long Haul Retransmission

There is no standard recommendation for reliable transmission of science data over deep space link for non-file data. LTP provides a transport-layer service through the use of a proven retransmission capability similar to the currently flying CFDP. CFDP provides a retransmission capability, but since CFDP is an application it is dependent upon the use of files.

F. Closing the Loop at the Ground Station

BP/LTP provides a method to “close the loop” at the ground station instead of at the mission control center. For Mars relay data, this would be useful in that the retransmission of corrupted relay data is requested at each hop. For instance, long-haul data from a relay agent can be retransmitted when it initially reaches Earth at the ground station instead of being transmitted to a JHU/APL MOC/SOC, routed to NASA/JPL or another organization’s MOC/SOC before retransmission can be requested.

A. Use of Thin Application Layers

Without a form a reliable retransmission in lower layers, applications are required to implement their own retransmission capability. As an example, CFDP has two methods for transmission of data: 1) acknowledged and 2) unacknowledged. Acknowledged mode guarantees delivery of data, but at a cost of extra complexity in the software. Additionally, this retransmission capability crosses over different layers in the protocol stack instead of relying on the underlying protocols (such as TCP in a terrestrial environment) to do the retransmission. BP/LTP allows CFDP to operate in a simplified, unacknowledged mode by moving the guaranteed delivery to a lower layer (LTP).

IV. Future Mars Network Architectures

As a step towards examining the use of BP/LTP in a Mars environment, the NASA/JPL and JHU/APL team developed two initial reference architectures for future Mars network needs and to provide a reference point for comparison of the BP/LTP scenario. The first scenario uses CFDP and the second BP/LTP. The nodes of the network consist of two landers (L_1, L_2), two orbiters (O_1, O_2), three Deep Space Network (DSN) stations (D_1, D_2, D_3), and one Mission Operations Center (MOC).

A. Using CFDP

Figure 2 shows a Mars network architecture that uses CFDP for data transmission. This scenario makes use of the CFDP Store-and-Forward Overlay (SFO) where each waypoint user (in this case, the orbiters and the DSN stations) takes complete custody of a file before it transmits elements of that file further towards its destination. In this example, data are collected by the Mars landers into files. When the bandwidth is available, the files are broken into smaller segments called Protocol Data Units (PDUs). These PDUs are transmitted to one of the two orbiters, O_1 or O_2 , along with Metadata PDUs which provide information about the file such as the filename. These PDUs flow over the CCSDS Proximity-1 protocol between a lander and an orbiter and the files are reconstituted onboard the orbiter⁷. Once the orbiter receives the complete file, it then performs another transfer of that file to the DSN ground stations over the CCSDS TM Space Data Link Protocol⁸. From there, the data are transferred to the appropriate mission data center via the Space Link Extension (SLE) services⁹. Throughout this process, the CFDP acknowledged mode is used to ensure delivery of the data through retransmissions of PDUs where necessary.

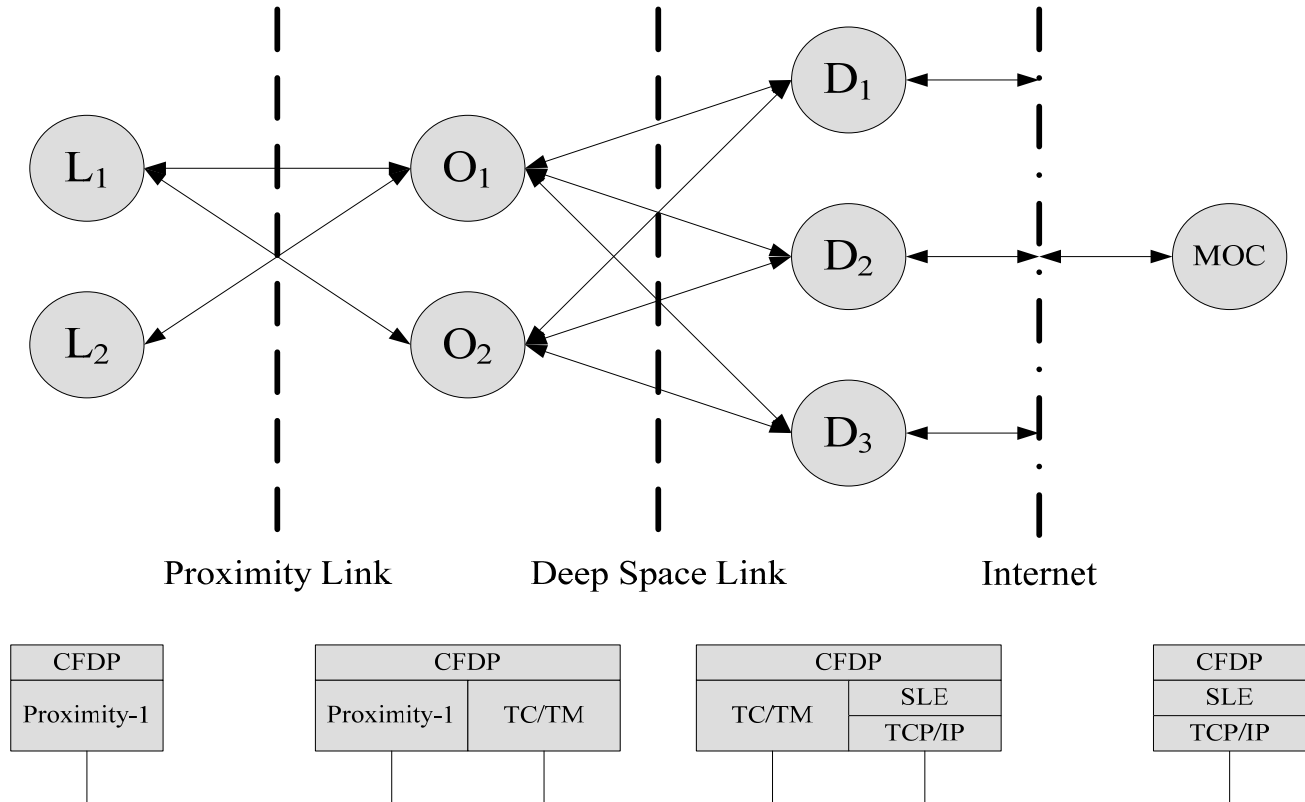


Figure 2. Mars Communication Architecture Using CFDP

A limitation in this architecture is that it requires full delivery of a file to an orbiter before that file can be transmitted to Earth. It does not allow for parts of a file to be transmitted to one orbiter and parts to another orbiter. As a result, files need to be tailored to fit within a given bandwidth margin or a given storage allocation. Also, if a pass is missed or an anomalous situation occurs with the orbiter, the transmission of the file's data will cease unless transmission of the file is manually rescheduled or a transaction may have to be cancelled midstream before a file can reach the ground station. The manual rescheduling of these links increases mission costs due to additional planning by the mission operations teams and delays the receipt of data by the science teams.

B. BP/LTP

An alternative solution to using CFDP SFO is shown in Fig. 3. CFDP is again used, but in this case unreliable mode is used where data is not automatically retransmitted by the CFDP protocol. Instead the CFDP implementation is simplified and the CFDP layer is made “thinner” by relying on underlying protocols (Proximity-1, LTP) to perform retransmissions. Specifically, CFDP provides directory access and file access capabilities. Also, instead of CFDP PDUs being the common unit across the network, the network uses BP bundles. To start, PDUs are placed in BP bundles and the bundles are reliably transmitted to the orbiter via the Proximity-1 protocol. Bundles are stored onboard the orbiter until bandwidth becomes available to transmit the data to Earth. At that point, the BP bundles are reliably transferred to DSN stations via LTP over TC/TM.

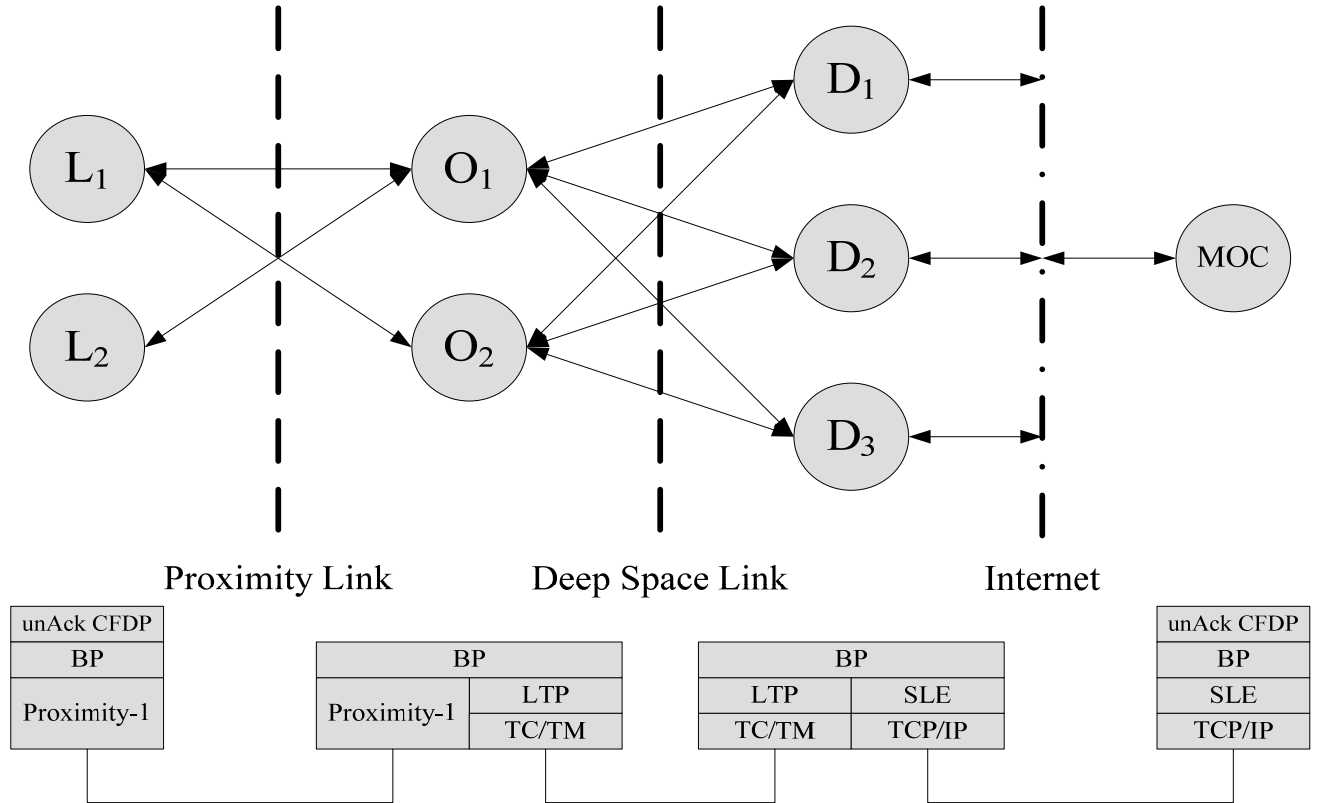


Figure 3. Mars Communication Architecture Using DTN Protocols

Unlike with CFDP SFO, BP/LTP allows for CFDP PDUs for a single file to be transmitted along multiple paths. If a bundle containing a PDU is not received by an orbiter or an expected orbiter path suddenly becomes unavailable, the bundles can be rerouted along a different path to the destination. Dynamic routing in DTN is a current research area. Since our main focus is not on developing dynamic routing algorithms, we assume complete knowledge of the dynamic changes of link availability in order to reroute bundles.

The protocols in our simulation are used to model communication in a potential Mars relay network. We are using the topology of and historical data pertaining to the current Mars Exploration Rover/Mars Odyssey/Mars Global Surveyor mission to observe the behavior and performance of these selected protocols. With the cancellation of the Mars Telecommunication Orbiter, science orbiters with similar orbital characteristics to Mars Odyssey and Mars Global Surveyor as well as Mars Reconnaissance Orbiter will perform relay operations. Therefore, the use of this data provides a high fidelity simulation of future relay orbiter characteristics.

V. DTN Mars Network Experiments

The team performed experiments in two different settings. First, we simulated the Mars network using the Multi-mission Advanced Communications Hybrid Environment for Test and Evaluation (MACHETE). Although

MACHETE can be integrated into a hybrid simulation involving software network simulator and testbed, in the first setting, we simulated the entire network in a discrete event network simulator, running faster than real-time. Secondly, we configured the Next Generation Mars Network Protocol Testbed at the NASA/JPL's Protocol Test Lab to validate our simulation by running live traffic.

A. Discrete Event Network Simulation

In this experiment, we used the MACHETE tool developed at NASA/JPL. The components of MACHETE are software modules. In this stand-alone experiment (not integrating MACHETE to testbeds), the simulation speed is much faster than real-time. In fact, we ran a seven day scenario in a few minutes.

1. Network Simulation Tool

To simulate network protocols, NASA/JPL developed a comprehensive tool MACHETE¹⁰, tailored to unique characteristics of space networks. The architecture for MACHETE consists of 1) orbital and planetary motion kinetics modeling, 2) link engineering modeling, 3) traffic load generation and space communications protocol models, and 4) external interfaces. At the core of MACHETE network simulator is a discrete event simulator QualNet developed by Scalable Networks, Inc. QualNet is the commercial product of GloMoSim which was developed as part of the DARPA Global Mobile communications networking project. QualNet contains a full contingent of conventional protocols such as the IEEE 802.11/WiFi and Internet protocol standards. The specific space protocol models developed at NASA/JPL are built upon QualNet; these include the complete CCSDS protocol stack: Proximity-1, Packet Telemetry/Telecommand, Advanced Orbiting System (AOS), Space Communications Protocol Standards (SCPS) and CFDP. The most recent additions are models for BP and LTP, built according to specifications of IETF drafts; the effort of BP and LTP modeling are reported in reference¹¹.

2. Mars Network Simulation

To simulate our scenario in MACHETE, we assume two generic science orbiters with a similar relay capability along with historical orbital representing expected pass times. We obtained a representative set of data covering a one-week period of relay pass data from the NASA/JPL Planning and Execution Systems Section. We set up our experiments using the corresponding rover-orbiter and orbiter-DSN passes and data rates. To simplify scheduling, we assume that only one orbiter can communicate to a DSN station at any time. This limitation will be removed in future simulations because the DSN is capable of supporting two concurrent communications each involving a unique orbiter and ground station pair.

Our simulator is capable of reading in the delay report files to associate the appropriate propagation delays to each link with respect to time. However, in this initial experiment, we use an average propagation delay computed from delay reports. The reason for this is to simplify the computation of the communication schedules. The simulator is also capable of reading in bit-error-rate profiles per link with respect to time. We did not incorporate bit-error profiles in the experiment because our focus is on the basic functionality of the protocols. We do intend to incorporate bit-error-rate profiles in future simulations.

We assume a data rate of 128 kbps between landers and rovers which is typical for the Mars Exploration Rovers. The link availability as well as data rates for the deep space links were provided by the Planning and Execution Systems team over the representative one-week period. The following tables summarize link availabilities, normalized to simulation period. As we can see, the landers have very limited opportunities to communicate with the orbiters but the orbiters have much more opportunities to communicate with the DSN.

Table 1. Lander-Orbiter Contacts

Contact (percent time)	Orbiter_1	Orbiter_2	Total
Lander_1	3.31%	3.44%	6.75%
Lander_2	3.24%	3.57%	6.81%

Table 2. Orbiter-DSN Contacts

Contact(percent time)	DSN_1	DSN_2	DSN_3	Total
Orbiter_1	1.72%	23.22%	15.2%	40.14%
Orbiter_2	17.69%	13.63%	24.98%	56.3%

All bundles require custody transfer. The Time-to-Live settings for bundles are set high enough to avoid bundle expiration. We assume buffers are large (unlimited) so that data is not dropped. Complete knowledge about link availability is assumed for routing. We do not use Linear Programming to optimize the schedule. Instead, we use a simple “First Contact”¹ and highest data-rate heuristic to route bundles. Although BP can handle data of different priorities, we limit the data traffic to be of the same priority for our first observations. In our experiment, we used the aforementioned network topology and ran BP over LTP over expedited Proximity-1 on lander-to-orbiter links and BP over LTP over TC/TM on deep space links. We assumed a 50% load (link utilization).

In the BP/LTP experiment, Lander_1 sent 277 bundles to the MOC and Lander_2 sent 280 bundles to the MOC. Each bundle is 1 Megabyte in size. The bundles relayed by Orbiter_1, Orbiter_2, and by DSN_1, DSN_2, and DSN_3 are shown in Fig. 4. All the bundles sent by Lander_1 and Lander_2 are received at the MOC, as expected. The propagation delay on a proximity link (between lander and orbiter) is 16 ms, the propagation delay on a deep space link is 552 seconds, and terrestrial links (from DSN to MOC) has a propagation delay of 1 ms. Note that the latency for a bundle from a lander to the MOC includes effects of propagation delay, link availability and schedule. The minimum, average and maximum latency for a bundle from a lander to MOC are 6 minutes, 2.4 hours

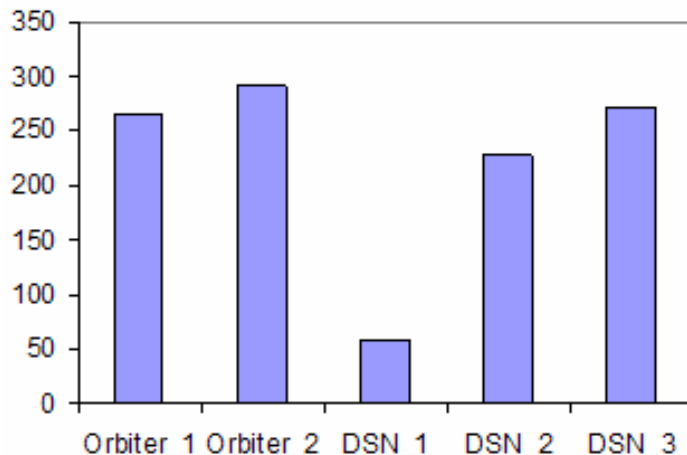


Figure 4. Count of Relayed Bundles.




Figure 5. Experimental Bundle Latency Distribution in Hours.

This exercise tested the basic functionality of the BP/LTP protocols in a multi-hop actual mission scenario with mission data running over typical deep space proximity-link communication protocols. We observed expected and acceptable performance of the protocols in terms of bundle latency and reliability. We conjecture that the latency can be improved by better scheduling algorithms than the simple “first contact” scheduling algorithm, but scheduling is outside the scope of this paper. This result shows that even in a real-world network topology running several complex protocols, BP and LTP proved robust.

and 20 hours respectively.

Distribution of the bundle delay (including propagation delay, link availability and schedule) is shown in Fig. 5, where the unit on the y-axis is hour. We observe that very few bundles have latency higher than 10 hours. We have correlated the experimental result with our analysis on average and maximum latency according to the specific schedule, and found that the experiment result matches our expectation. Figure 6 shows the predicted latency distribution based on the schedule but ignoring queuing and transmission delay. The predicted and resulting distributions are reasonably similar.




Figure 6. Predicted Bundle Latency Distribution in Hours.

B. Hybrid Mars Network Simulation

The hybrid experiment includes a combination of simulated protocol stacks and the NASA/JPL reference implementation of CFDP. The setup also includes the Channel Simulator to create realistic deep space link conditions. This setup makes an extensive use of NASA/JPL's Protocol Test Lab capabilities as described below.

1. Protocol Test Lab (PTL)

The Protocol Test Lab's mission is to provide an end-to-end test environment for performance benchmarking and validating the functionality and interoperability of various deep space communication protocols. To mimic the deep space communications architecture, PTL employs a combination of simulation and emulation tools to create realistic space link behavior such as delay and frame error rates along with a combination of protocol stacks. The PTL is equipped with a cluster of thirty-five rack-mounted computers, managed switches, a VPN concentrator, RAD6000, RAD750 flight processor hardware and chassis and UHF radios to provide a rich set of space communications protocol simulation options. Figure 7 summarizes the functional components and their usage at PTL.

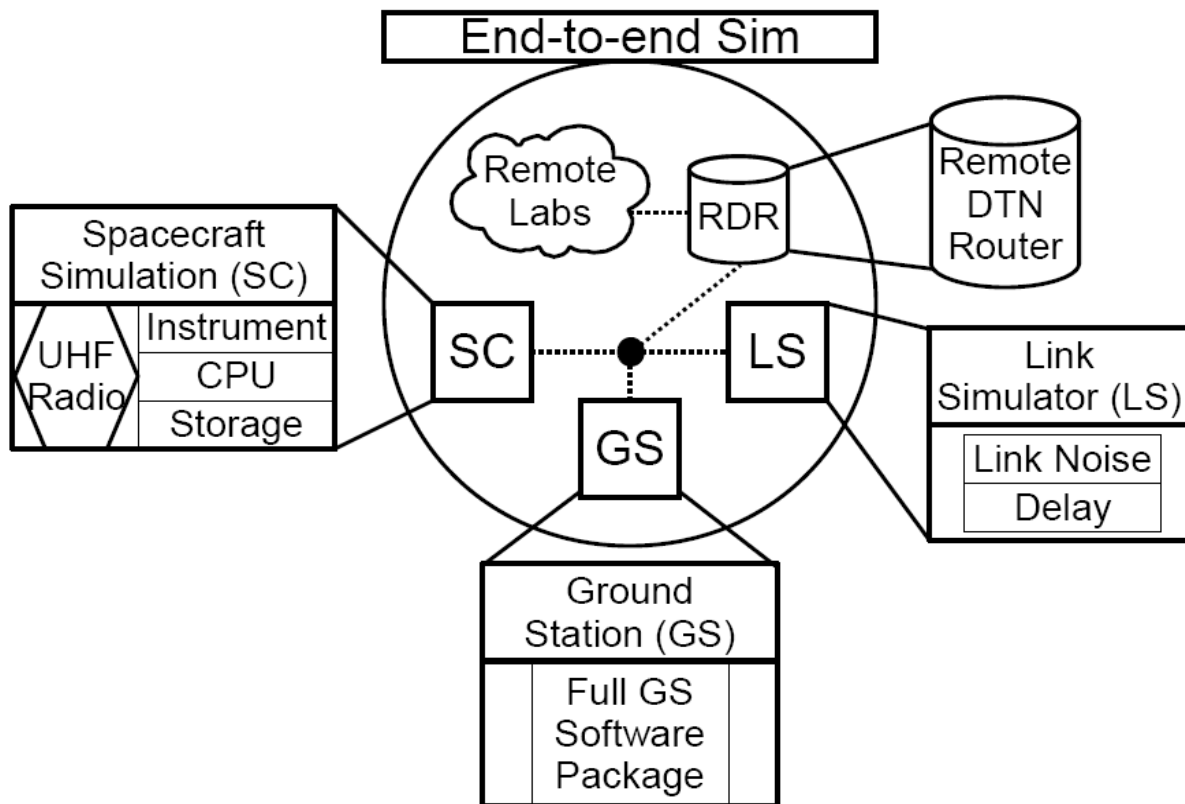


Figure 7. Protocol Test Lab Functional Diagram

2. Testbed Configuration

As shown in Fig. 8, the hybrid testbed is made up for four distinct components, namely the NASA/JPL reference CFDP running over UDP and TM/TC link layers, the Channel Simulator that simulates deep space link conditions and the MACHETE node that simulates Proximity-1 link between the landed assets and the Orbiters. MACHETE, when integrated into the testbed, runs in real-time mode.

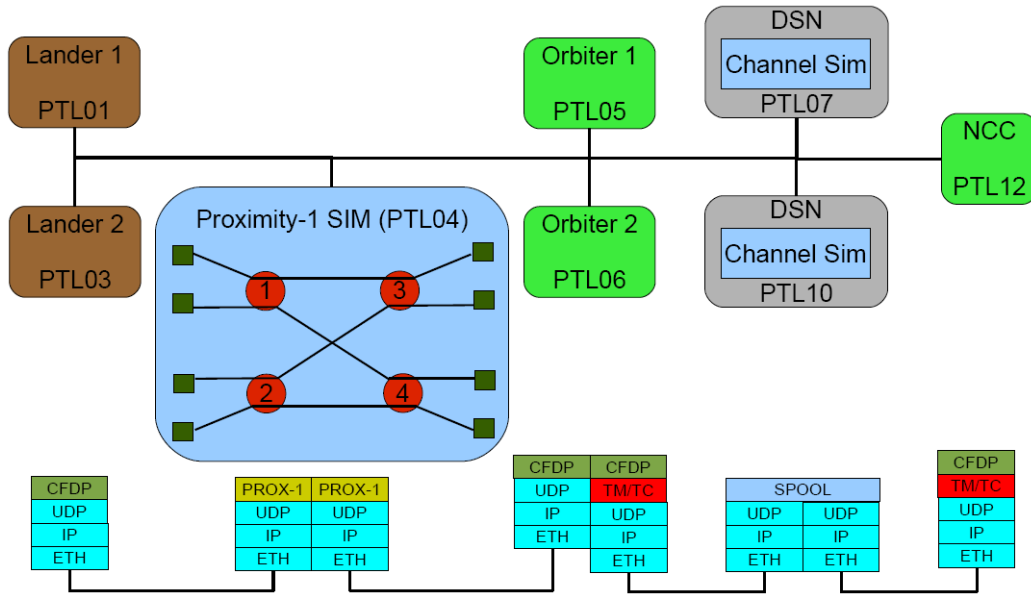


Figure 8. Protocol Test Lab Mars Network Testbed

The landers, orbiters and the Network Control Center (NCC) ground station each runs CFDP with Extended Procedures to support CFDP relaying across multiple nodes. We run CFDP over UDP between the landers and the orbiters to allow the virtual node in MACHETE (representing the sending node) to capture CFDP messages and pass the information to its Proximity-1 simulated protocol model so that communication effects occur. The Channel Simulator module has the ability to change the link conditions based either on the static frame error rate or the fluctuating link condition models outputted from the Satellite Orbital Analysis Program (SOAP, developed by the Aerospace Corporation). Between the Orbiters and NCC, CFDP is run over TM/TC to create the long haul CFDP protocol stack structure. The whole testbed is running on 100 Mbits/s Ethernet network isolated from all other network traffic except the traffic related to this experiment.

Our preliminary results show that we have the capability to run this scenario using the Suspend and Resume feature of CFDP to follow the Mars occultation schedule. These initial results show that using CFDP with Suspend and Resume feature can save manual scheduling operational cost and increases the amount of data returned to earth compared to transmitting science data without using CFDP. The detailed quantitative analysis of the increased data return rate and operational cost saving is premature to report in this paper.

VI. Conclusion and Future Direction

As future missions will involve increasingly complex systems, it becomes more important to have standardized communications among the entities. Currently, there is already a suite of CCSDS protocols that support space networking. In this study, we investigate using DTN in conjunction with these standards as well as terrestrial protocols as a common overlay for network entities that may run differing transport protocols.

Using a discrete event network simulation tool and NASA/JPL's models of BP and LTP, we ran experiments to observe the performance of BP over Proximity-1 or LTP. Actual historical mission data on link availability and data rates were used by the simulator. We used a simple scheduling algorithm to compute communication schedule. From our simulation, we observed BP showed almost no performance impact in an actual mission scenario. BP has the additional advantage of being a common overlay that allows for more flexibility in protocol combinations of lower layers.

In addition to using mission data, we built a model in the Satellite Orbital Analysis Program that produced link availability and propagation delays that matched the mission data. This will enable us to generate contacts for time periods other than the one provided by the mission planning team. Future extensions will include stress testing the protocols with higher traffic loads and studying the usage of buffer space. Another extension to this work would be to examine the adaptability of BP to unexpected link/node outages by re-routing dynamically. One approach is to incorporate known research results on DTN routing into the computation and updates of communication scheduling and routing.

Acknowledgments

The authors would like to thank Bruce McLaughlin, Grailing Jones and Paul Fieseler from JPL's Planning and Execution Systems Section for their input on representative Mars exploration mission data, spacecraft attributes, and assistance with correlating mission data with simulation data. The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology and The Johns Hopkins University Applied Physics Laboratory, under a contract with the National Aeronautics and Space Administration.

References

- ¹Jain, S., Fall, K., and Patra, R., "Routing in a Delay Tolerant Network," *ACM SIGCOMM*, Aug. 30 – Sept. 3, 2004, Portland, Oregon.
- ²Burleigh, S., Fall, K., Cerf, V., Durst, R., Scott, K., Weiss, H., Torgerson, L., and Hooke, A., "Delay Tolerant Networking – An Approach to Interplanetary Internet," *IEEE Communications Magazine*, June 2003
- ³Burleigh, S. et al., "Licklider Transmission Protocol" (main protocol), draft-irtf-dtnrg-ltp-03.txt, July 2005.
- ⁴CCSDS 727.0-B-3. "CCSDS File Delivery Protocol (CFDP)". Blue Book. Issue 3. June 2005.
- ⁵Scott, K., Burleigh, S., "Bundle Protocol Specification", draft-irtf-dtnrg-arch-02.txt, July 2004.
- ⁶Cerf, V., Burleigh, S., Hooke, A., Toergerson, L., Durst, R., Fall, K., and Weiss, H., "Delay-Tolerant Network Architecture," INTERNET-DRAFT, July 2005.
- ⁷CCSDS 211.0-B-3. "Proximity-1 Space Link Protocol—Data Link Layer". Blue Book. Issue 1. September 2003.
- ⁸CCSDS 132.0-B-1. "TM Space Data Link Protocol". Blue Book. September 2003.
- ⁹CCSDS 910.0-Y-2. "Space Link Extension Services—Executive Summary". Yellow Book. Issue 2. March 2006.
- ¹⁰Gao, J., Jennings E., Clare, L. SeGui, J., Kwong, W., "MACHETE: A Tool for Architectural Modeling, Performance Characterization, and Technology Infusion of Space-Based Networks," AIAA International Communications Satellite Systems Conference (ICSSC), Rome, Italy, September 2005.
- ¹¹Seguí, J., Jennings E., "Delay Tolerant Network Protocol Simulation in MACHETE", (to appear), 2nd International Conference on Space Mission Challenges for Information Technology (SMC-IT), Pasadena, CA, July 2006.