Moving Toward Space Internetworking via DTN: Its Operational Challenges, Benefits, and Management

Erik Barkley[1], Scott Burleigh[2], Roy Gladden[3], Shan Malhotra[4], Peter Shames[5].
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

The international space community has begun to recognize that the established model for management of communications with spacecraft – commanded data transmission over individual pair-wise contacts – is operationally unwieldy and will not scale in support of increasingly complex and sophisticated missions such as NASA's Constellation project. Accordingly, the international Inter-Agency Operations Advisory Group (IOAG) [i]chartered a Space Internetworking Strategy Group (SISG), which released its initial recommendations in a November 2008 report. The report includes a recommendation that the space flight community adopt Delay-Tolerant Networking (DTN) to address the problem of interoperability and communication scaling, especially in mission environments where there are multiple spacecraft operating in concert. This paper explores some of the issues that must be addressed in implementing, deploying, and operating DTN as part of a multi-mission, multi-agency space internetwork as well as benefits and future operational scenarios afforded by DTN-based space internetworking.

I. Introduction

Space missions have typically been designed to operate via direct communications to and from Earth. While this configuration is still relevant for certain missions, for some time this "stand-alone" approach has been changing, driven by a variety of factors. Primary among them are the benefits derived from collaborative missions that would at least share data and that may operate interdependently to increase the overall science return. A prime example of this is relaying data from smaller, lower powered (thus less costly) planetary landers using a nearby, more capable orbiter. In some cases there may be benefit in having one spacecraft directly signal another, such as in notification of natural events (tsunami, tornado), or in passing data from one remote spacecraft or outpost to another.

Data relaying occurs today, but the relay capabilities have developed over time, often as add-ons to missions whose nominal task was science and whose communications architecture and operations model was also focused on doing that mission's science. This is not an optimal approach if an additional mission objective is to provide network or relay services to other missions. For the latter to succeed in any broad sense, all of the cooperating missions must adopt some common and interoperable protocols and processes that will actively support the delivery of these networked services, and do so in a fashion that is reliable, consistent, easy to operate, and extensible.

This paper will:

- Briefly survey the drivers for a space internetworking approach and the international agreements that are strongly motivating its development.
- Identify the primary benefits that result from implementing this approach via DTN.
- Survey the current NASA-JPL communications infrastructure, including its Deep Space and Mars Relay Networks.
- Summarize DTN proof-of-concept efforts conducted with an operating deep space mission.
- Identify the key technical challenges and operational considerations in moving towards ubiquitous DTN-based space internetworking.

The paper will conclude with a discussion of some future application scenarios that are enabled via this approach.

[1] Principal Engineer, M/S T1716, 4800 Oak Grove Drive, Pasadena, California, 91109
[2] Principal Engineer, M/S 301-490, 4800 Oak Grove Drive, Pasadena, California, 91109
[3] Multimission Relay Operations Lead, Mars Network Office, M/S 264-2354800 Oak Grove Drive, Pasadena, California, 91109
[4] System Engineer, NASA Deep Space Network, MS 144-200 4800 Oak Grove Drive, Pasadena, California, 91109
[5] Manager, JPL Data System Standards, M/S 301-490, 4800 Oak Grove Drive, Pasadena, California, 91109

Copyright 2010 California Institute of Technology. Government sponsorship acknowledged. Copyright 2010. All rights reserved.

II. Brief Survey of International Developments

The major international civilian space agencies formed the Interagency Operations Advisory Group (IOAG), in June 1999; it meets on a yearly basis. The IOAG fostered wide adoption of a standardized, link-layer, cross-support architecture, the Space Link Extension (SLE). In recognition of mission plans being formulated by multiple agencies for future interdependent mission sets, particularly at Mars and the Moon, and prompted by a Communiqué from its parent organization, the international Inter-Operability Plenary (IOP), the IOAG created a Space Internetworking Strategy Group (SISG). The SISG is chartered to "… formalize a draft Solar System Internetwork (SSI) Operations Concept and candidate architectural definition … and should prepare a mature architectural proposal for review and endorsement at the third Inter-Operability Plenary meeting (IOP-3)."

The following IOAG resolution constituted the charter for the SISG: "The IOAG resolves to form a Space Internetworking Strategy Group to reach international consensus on a recommended approach for transitioning the participating agencies towards a future 'network centric' era of space mission operations. The group will focus on the extension of internetworked services across the Solar System, including multi-hop data transfer to and from remote space locations and local networked data interchange within and among the space end systems."

The IOAG/SISG has published a report[ii] which contains an extensive set of findings and recommendations, some of which are provided here for context:

> The SISG has reached a consensus recommendation that the international community should begin the planning and development activities that are necessary to transition future space mission operations to rely on a new end-to-end internetworked model of data communications, using mission support infrastructure that spans across space. To achieve this goal, the SISG therefore recommends that the participants in the second Inter-Operability Plenary (IOP-2) in December 2008 should be asked to approve a significant international initiative that will begin the build-up of required infrastructure which will evolve towards a proposed "Solar System Internetwork" (SSI) in support of future missions of space exploration. The SSI is envisaged to be a voluntary confederation of space and ground based infrastructure that is contributed by individual participants (including government and commercial entities) and that is bound together by a common architecture and the adoption of common standards.

Selected top level SISG Summary Conclusions include:

- **SC-01**: The SISG notes that a major factor in enhancing the overall productivity of many space missions is their ability to draw upon shared communications and navigation services that are enabled by international interoperability and cross-support…
- **SC-03**: Building upon this current success of terrestrial cross support, the SISG concludes that the international community now needs to address the development of future in-space cross support services (including cross-support on and around other Solar System bodies). The current in-space cross-support record is spotty …In order to realize this end-to-end cross-support architecture (which is expected to become a key enabler for new international space mission initiatives), the SISG also concludes that space-faring entities should be encouraged to begin extending their cross-support services into space via the build-up of re-usable, confederated in-space communications and navigation infrastructure which can offer services across organizational boundaries…
- **SC-04**: Analysis of emerging Earth observation, Lunar and Mars exploration mission scenarios indicates that an increasing number of individual space missions will need the ability to confederate and share in-space communications resources and infrastructure in order to achieve goals that are greater than any one of them could individually accomplish…
- **SC-05**: The choice of space internetworking protocols is important. While there are regions of space in which conventional IP-based networking is expected to work well, it is inherent in the physics of space communications that between (and at the edges of) local regions of short-delay "always on" IP connectivity there will be many areas of disrupted or long-delay communications."

The SISG notes that the emerging Disruption Tolerant Networking (DTN) technology is the only mature candidate protocol available that can handle the disconnection and delays inherent in many regions of space operations. Therefore the SISG concludes that future space networks should provide DTN as a primary end-to-end routing service and that the development of DTN to full flight readiness circa 2012 should be a high priority. To meet projected mission requirements, the early phase of the SSI should begin in the 2015 timeframe to support launching missions and by about 2025 a routine operational capability should be in place.

Since that report was published the work of the SISG has continued, developing a set of requirements, a Concept of Operations, and an initial draft of an Architecture Definition Document. In support of the IOAG and SISG resolutions and to support future mission sets at the Moon and Mars, NASA has embarked on an aggressive program to standardize and mature the DTN protocol suite.

III. Benefits of DTN Implementation

Traditional "non-networked" communications architectures require that all communication to/from a user spacecraft be brought to/from Earth on direct, point-to-point links. For example, if two spacecraft in the Lunar or Mars environment wish to communicate, whether they are in orbit or on the surface, they require a line-of-sight communication with Earth, which introduces a "two hop" path from the first spacecraft to Earth and then back to the second spacecraft. This introduces unnecessary light time delays into the communication channel. In contrast, if it is possible to support local, in-situ routed networking, it becomes possible to confine local communications to the proximate environment, without the long round-trip to Earth.

For missions at Mars, where the communications delays are far greater than those in the near-Earth environment, the benefits of these relayed operations are significant. For a number of mission designs, relaying offers significantly higher data returns, at lower power and with less mass, than can be achieved with direct-to-Earth approaches. Many of these benefits are available to missions that utilize simple file relaying operations, but with the adoption of DTN further benefits are possible in terms of interoperability, more straightforward cross support, and operational simplicity. As noted in the IOAG report, current cross support is "spotty" and complicated largely because there are not yet in place well-tested and broadly-adopted standards.

Current approaches, while they do adequately function, amount to a set of ad hoc, "bucket brigade" transfers of data, each of which uses locally defined protocols and procedures. These approaches are human intensive, require detailed knowledge of each possible relay spacecraft and different implementations for each, are hard to coordinate, and do not scale well. DTN, on the other hand, is designed to provide a space internetworking layer similar to the familiar terrestrial Internet, but tailored for operation in space where delays and sporadic connectivity are the norm. Under this layer the familiar space links, with their coding gain advantages, will still operate. On top of these DTN links existing file delivery standards (e.g., CFDP), new message delivery standards (e.g., AMS), and other public or even mission-unique standards can be deployed. In a multi-mission environment these each provide end user functionality, but also inherit the benefit of reduced implementation complexity because of the standardized routing, planning, and store-and-forward services provided by the DTN.

IV. Current JPL Infrastructure

To facilitate an understanding of current NASA/JPL data communications and help provide a basis for appreciating DTN-based space internetworking, it helps to examine the current deep space infrastructure, composed primarily of the Deep Space and Mars Relay Networks.

A. Deep Space Network (DSN)

The Deep Space Network (DSN) provides communication and tracking services to and from various spacecraft around the Solar System. It does this via more than 16 Earth-based antennas that are coordinated to track the various spacecraft. Considering the Earth as a rather large spacecraft with many antennas, it is possible to think of existing DSN/spacecraft interactions as a precursor to DTN operations, especially with regard to the management of individual pair-wise contacts.

In contrast to typical terrestrial Data Link Layers, the space link for deep space missions requires significant information and processing facilities for its management. This involves trajectory data, mutual communications view/geometry calculations, mutually agreed space link configurations (coding, modulation, link protocols), space link service schedules, space link Doppler compensation, and space link production adjustments in regard to spacecraft induced artifacts.

Trajectory information is submitted to the DSN for each spacecraft that is to be tracked. From this information, communications geometry is calculated, along with other physical antenna-unique events of interest, such as the clearance of various elevation thresholds and spacecraft/planetary occultations. This information yields what is physically possible in terms of communication with various spacecraft.

Based on this set of physically feasible DSN communications opportunities, mission operators determine which opportunities are to be utilized and scheduled with the DSN[6]. The DSN subsequently releases this schedule information to all of the missions to which it provides service. This is generally negotiated well in advance of the actual communications sessions, with a general goal of a very stable schedule being available at least eight weeks beforehand. However, it should be noted that there are often adjustments and re-releases of the schedule within the eight week timeframe due to various changes in mission and/or DSN circumstances.

[6] Scheduling DSN services with regard to many mission spacecraft is considerably more involved than described here. However, the details of that complexity do not substantially alter the concepts presented in this paper.

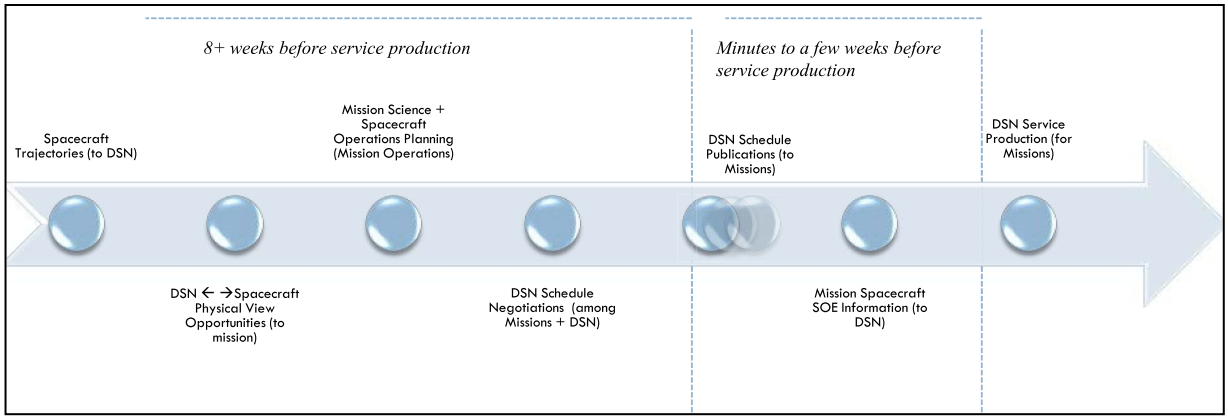

Figure 1. DSN Space Link Management Overview.

In addition to the physical information necessary to coordinate communication and tracking services, logical information such as general spacecraft availability for communication (e.g., in addition to being in mutual view, is the spacecraft's antenna pointed at Earth?) and the state of that communication's availability (e.g., selection among multiple possible bit rates) also needs to be considered. This is typically conveyed to the DSN in a time-ordered report which describes various changes in the space link due to spacecraft considerations such as changes in bit rate, whether or not telemetry is being modulated, etc. It should be noted that this time-ordered information must, by definition, be in conformance with the physical view/communications geometry information and scheduled service times. This information is typically supplied to the DSN two to four weeks before actual communications sessions, but similar information is sometimes communicated just prior to or during a session.

Physical space link management is achieved by mathematically modeling each pair-wise contact such that the communications geometry-induced Doppler dynamics are kept to an absolute minimum from the spacecraft antenna perspective. Shortly before the time of communications, the DSN internal systems are configured with, amongst other configurations, the Doppler "curves" to enable a "smooth" frequency signal for the mission spacecraft.

The overall DSN space link management process is summarized in Fig. 1. Today, portions of this process are somewhat automated, but there is still considerable room for improvement. Tools are currently in development to facilitate and automate portions of the scheduling process. A framework is being established to support automated, on-demand communication geometry computations, and a scheduling engine is in development to allow users to generate schedules of space link services in a more automated, rapid fashion. In addition, a comprehensive data collection and analysis effort is currently underway that will improve reporting on DSN performance and service quality.

B. Mars Relay Network

During the middle of 2008, there were six spacecraft at Mars that could communicate with each other in a "relay network," as in Fig. 2. The twin Mars Exploration Rovers (MERs), Spirit and Opportunity, arrived at Mars in January 2004, and they have been relaying data to Earth via the Mars 2001 Odyssey orbiter since that time. The European Space Agency's Mars Express Orbiter has been in orbit at Mars since just before the arrival of the rovers and has performed several relay tests with them, but has not served as a primary relay asset. The Mars Reconnaissance Orbiter (MRO) arrived at Mars in 2005 and has

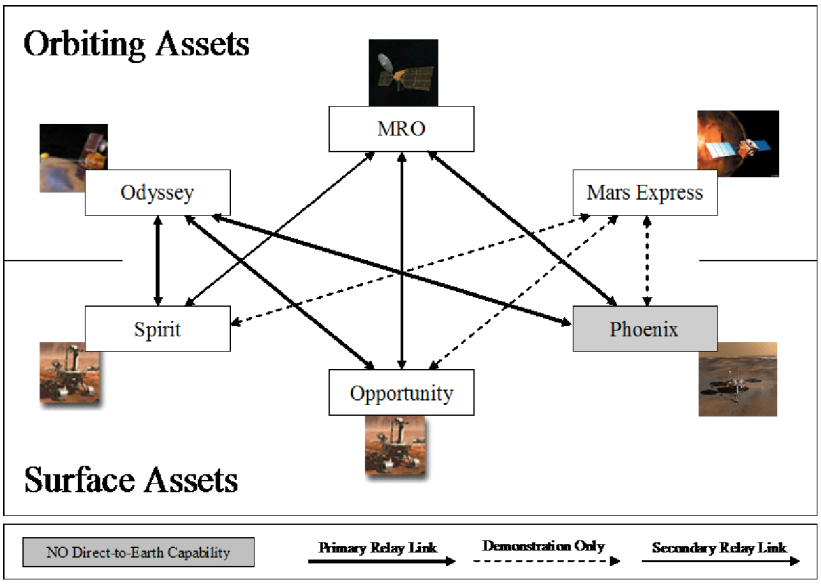

Figure 1. The 2008 Mars Relay Network.

concluded its primary science mission, but it has not to-date been providing primary relay support to either of the rovers. Most recently, the Phoenix lander arrived in May of 2008 and utilized both Odyssey and MRO for primary relay operations during its five month mission.

This evolving and ad hoc network of spacecraft was composed of vehicles that, in some cases, weren't designed to smoothly interoperate and interface with each other. For example, within this small network three different telecommunications protocols were implemented to transfer data from one vehicle to another. This reflected the changing nature of the network itself, as well as advances in telecommunications and other spacecraft technologies. However, since each spacecraft handled relay data in different fashions, the operators of the landed spacecraft were required to have detailed knowledge of the orbiting spacecraft that provided relay services.

With the exception of the Phoenix lander, relay communications as part of the Martian relay network was not implemented by the various spacecraft projects as a primary or key driver to accomplish mission objectives. For the orbiters, which had their own scientific goals to pursue, providing relay support to the landed vehicles interfered with ongoing mission objectives and relay communications were managed as a secondary priority; the rovers were designed to achieve mission success in the absence of relay support.

In addition, each vehicle in the network was designed and managed independently and the operators of these spacecraft had differing operational processes, procedures, planning timelines, and – not to be discounted – culture. Extensive negotiations occurred to identify an operational approach that would facilitate the transfer of data between the various spacecraft in spite of these operational disparities. In most cases, software and ground data system interfaces were originally designed to meet primary mission objectives, and relay activities were treated as an add-on capability. As a result, these interfaces had to be carefully negotiated as did the human interfaces.

Ultimately, the interfaces were developed to provide a functional, if idiosyncratic, infrastructure that could be used to identify windows of opportunity to transfer data between the various spacecraft and to implement transfers during those windows. The relative geometries of the spacecraft, as determined by the orbiters' estimated position as they rose and fell from the perspective of the landers, provided the initial estimates for when these communication sessions could occur. From that initial plan further constraints were overlaid and negotiated between the spacecraft operators until a final communications plan was agreed upon. This negotiation effort was a largely manual process, with only minimal automation included, and is analogous to the previous discussion about how DSN services are scheduled and managed; it can be thought of as the Data Link management component of the Mars Relay Network.

At the time of this writing, new software is being developed to standardize and centralize this planning and coordination data. The intent is to provide a common interface for use by all the lander and orbiter projects that participate in the network so that the quirks of each of the various spacecraft can be largely hidden from the community. By centralizing the coordination data, this new software represents the first step towards further automation to simplify and streamline the transfer of data through the network.

However, even with the inclusion of this new software, there are fundamental features of each of the spacecraft in the network which prevent a transition to full DTN-enabled networking. For example, each of the orbiters manages the data to be transmitted to a lander as a "lump" of binary data. The orbiters do not process nor manipulate the data in any way to understand the destination of the data or to verify data integrity, but instead rely upon external meta-data, file names, and, in some cases, the time of receipt of the data to determine when and to which lander the data should be transferred.

To forward this data through an orbiter, the lander operators must first transfer it to the orbiter mission operations center where the data is managed as if it were an orbiter data product. The process to handle and transmit the data product to the orbiter from ground-based stations requires human involvement (though it could be automated), and there is nothing in the data product itself, save for the name of the file, that distinguishes it as anything other than an orbiter data product. Indeed, once the data is received onboard the orbiter, it continues to be treated as orbiter data until just prior to the scheduled relay session, when the data is finally handed off to the orbiter's radio transceiver and transmitted to the lander.

The transfer of data from the lander ground data system to the orbiter ground data system can be performed at any time, but the transmission from there to the ground station and then to the orbiter requires human involvement and is carefully planned and coordinated. Similarly, the transfer of the data from the orbiter to its final destination on the lander is carefully managed and requires unique and specific commanding on the part of the orbiter to enable the transfer of the data. None of the vehicles in the network are constantly receptive to the receipt of data or capable of transmitting data at any time; all of this must be carefully choreographed and specifically commanded.

Because of the human involvement in these processes, the technical interfaces onboard the orbiters, and the amount of operational overhead required to plan, coordinate, and implement communications sessions end-to-end, the current network by its very nature does not support "custody transfer." The human operators are keenly aware

of the details of the path by which data is both forwarded to and returned from the landers and are involved at nearly every point.

Efforts to automate several of the steps in these end-to-end communications chains have been met with some resistance by the various projects. For example, the lander operators are typically anxious to know of the latest time at which they can submit a data product to the orbiter operations team with a high degree of confidence that the data product can be transferred to the orbiter in time for a designated communications session with the lander; this is in order to maximize the amount of time that the lander operators have to generate said product. This issue is one of understanding the latest time at which data can be inserted into the network so that it arrives on time, which is a distinctly different issue from understanding how quickly data reaches its destination once inserted into that network. This latter point tends to be a key consideration when addressing network behavior, as most terrestrial network applications are generally more concerned about minimizing end-to-end transfer time, not about predicting the latest time of network insertion relative to an expected delivery time.

Because of this concern, the lander operators typically resist efforts to automate the chain of delivery of the data products because they want to personally know when the data moves from one point in the network to the next and be assured that some person is ensuring successful transfers and is able to intervene at the first sign of an issue. This particular concern could be addressed by the successful demonstration of a highly reliable and verbose system, but such a system is believed to be costly to implement, especially given the current ad hoc approach to communications protocols and architecture, and current network participants have neither the time nor the money to invest therein.

To complicate matters further, the current radios, as implemented on the orbiters, must be directly commanded to initiate a communications session; the orbiter radio is not "always on" but must rather be the hailing agent which establishes the communications session with the lander, meaning that the lander cannot itself initiate a session. In addition, the radios are not independent onboard the flight systems and interfere in some cases with the onboard scientific instruments due to electromagnetic interference (EMI), power, or thermal concerns; also they share the same data storage devices and on-board data buses as the remainder of the vehicle. Therefore, relay sessions represent a direct trade-off between the amount of science data that can be acquired and the amount of relay data that can be transferred to and on behalf of a lander.

All of these issues, however, could be mitigated in the near future on new vehicles by the implementation of a few simple modifications:

- The radio should operate independently of the remainder of the flight system without regards to power, thermal, or EMI constraints. It should have direct access to the communications systems that connect it to the network. Ideally the radio itself should have sufficient storage to manage data being transferred in either direction.
- The radio should be capable of being "always on", such that it is constantly prepared to transmit data to a lander at any time and receptive to a communications session initiated from another vehicle. This would aid in spacecraft emergency and other related commanding scenarios which the current network is ill-equipped to manage, and it would provide a mechanism to de-couple the commanding of the radio from that of the rest of the spacecraft.
- The radio should implement not just link layer protocols, with automated retransmission enabled, but the necessary protocols to implement the DTN suite and associated file or message delivery protocols.
- Custody transfer strategies should be implemented so that the onboard radio can receive and act upon data products as they are received either from a lander or from Earth, with full understanding of the source and destination of the product and the priority with which it should be transferred.

These modifications do not represent an exhaustive list of changes that would be required, but introducing them would partially enable DTN in the Martian relay network.

C. Deep Impact Network (DINET) Experiment

Significant work has been completed in demonstrating proof-of-concept for DTN-based space internetworking. Implementations of the DTN protocols and supporting infrastructure were operated in interplanetary space on the EPOXI spacecraft during the Deep Impact Network (DINET) experiment, from October 18 to November 13 of 2008[iii].

The core capability exercised by DINET was ION (Interplanetary Overlay Network), a multi-mission, multi-platform, reusable implementation of DTN protocols. The ION protocol implementations were integrated into both EPOXI's VxWorks-based flight software and JPL's Solaris-based ground data system, and they were additionally exercised in stand-alone mode on Linux workstations in the DINET Experiment Operations Center (EOC).

ION includes highly configurable, reusable implementations of:

- Bundle Protocol (RFC 5050[iv]), including Contact Graph Routing, Compressed Bundle Header Encoding, and Extended Class of Service
- Bundle Authentication, as defined by the Bundle Security Protocol (BSP) Internet Draft
- Licklider Transmission Protocol (LTP, RFC 5326[v])
- CCSDS File Delivery Protocol (CFDP, CCSDS 727.0-B-4[vi])
- CCSDS Asynchronous Message Protocol (AMS, CCSDS 735.0-B-1[vii])

These protocol implementations, which can be run on Linux, FreeBSD, OS/X, Solaris, VxWorks, or RTEMS platforms without modification, support the following elements of DTN functionality:

- Automated forwarding: received bundles are automatically queued for transmission on the best computed route toward the destination and are automatically de-queued when an opportunity to transmit on that route arises.
- Dynamic routing: different routes to a given final destination may be computed for different bundles, depending on requested quality of service, transmission backlog, anticipated transmission opportunities, and detected transmission anomalies.
- Priority: higher-priority bundles are transmitted before lower-priority bundles, not just at the source but also at every forwarding point.
- Flow and congestion control: rates of data transmission are automatically controlled at all nodes to prevent storage resource exhaustion. Bundle refusal may result in the computation of alternate routes.
- Delay-tolerant retransmission: reliability is assured by selective positive and negative acknowledgment and timeout-initiated retransmission that automatically adjusts for large variations in round-trip time.
- Custody transfer: upon successful bundle reception a node may "take custody" of the bundle, ensuring that any subsequent transmission failure will result in retransmission by that custodian node rather than the original source of the data.
- Delay-tolerant file transfer with incremental record delivery, enabling different portions of a file to be reliably forwarded in parallel through multiple forwarding nodes (orbiters and/or ground stations).
- Delay-tolerant, reliable publication or private transmission of brief messages, such as telemetry samples or reconfiguration directives.

In addition, a network management capability based on a DTN-enabled "load and go" mechanism implements a comprehensive ION network management system, which enables experiment network administrators to propagate changes in contact plans (e.g., schedules of available space link communications sessions, as discussed earlier), revise spacecraft clock correction deltas, stop and restart convergence-layer protocol adapters, add or remove Bundle Protocol endpoints, revise security policy rules, etc., anywhere in the network. All nodes automatically notify the network administration center of all processing anomalies and significant state changes, in addition to periodically returning summary network activity statistics; a network management console developed for the DINET EOC manages the display and storage of this information.

Since the conclusion of the DINET experiment, the ION software has been additionally deployed at several other NASA sites: some NASA centers are basing their participation in NASA's experimental DTN Engineering Network (DEN) on ION, and there are currently two ION-based DTN nodes resident on the International Space Station (ISS).

Integration of automated network protocols into a ground data system typically entails some rethinking of mission communications architecture. When the only information radiated toward a spacecraft is command data that may affect spacecraft operations, it is entirely appropriate for the mission operations center to subject all such information to close review and to release it only on specific authorization. But where the scope of uplink transmission is expanded to include data units of protocols that can only enact automatic communications and network activity – file reconstruction, delivery of messages to instruments, routing through relay points – those review and authorization procedures are no longer necessary and can fatally degrade the operation of the protocols.

DINET was able to use DTN protocols efficiently in ground data software by leveraging off the automated telecommand frame forwarding that had been developed to support Deep Impact's use of CFDP. As DTN protocols are infused into the core ground data system, support for automated forwarding of uplink frames containing native DTN data will need to be developed. The use of different virtual channels – either at the mission operations center or at the DSN tracking stations themselves – to segregate DTN uplink traffic from spacecraft commands can preserve necessary command review and authorization measures while still accommodating DTN's need for expeditious frame forwarding.

Several other infrastructure enhancements are also needed in order to fully integrate Delay-Tolerant Networking into mission operations at JPL:

- Automated forwarding of downlink frames containing native DTN protocol data units will be required. Virtual channel segregation, either at the mission operations center or at the tracking stations, might make this relatively simple.
- Standardized and well-documented mechanisms by which DTN-enabled CFDP and AMS functionality is offered to users must be developed.
- Further DTN security measures, as defined by the BSP Internet Draft, must be implemented in order to ensure the safe utilization of DTN infrastructure by a widely distributed JPL mission community. In particular, the bundle Payload Integrity and Payload Confidentiality protocols will be needed by science teams to assure the authenticity and reliability of instrument telemetry while preventing inadvertent or premature disclosure of valuable science data.

V. Key Technical Challenges

DTN's automatic computation of efficient forwarding routes and its automatic utilization of transmission and reception opportunities both rely on accurate anticipation of planned radio contacts among communicating entities (spacecraft and ground tracking stations). Enabling that anticipation is a two-part problem.

First, planning radio contacts has always been an operational challenge, both to the missions and to the Deep Space Network. Automated methods for using orbital geometry information to identify potential contacts have becoming increasingly powerful. Together with improvements in collaboration technology that simplify the reconciliation of planned operational activities with communication opportunities, these methods will make contact planning easier in coming decades.

Second, given a contact plan in a standardized digital representation, the remaining challenge is the propagation of that plan to the network nodes that will need to consult it. The network itself can normally be used to accomplish this network management task, provided the plan doesn't change after the last previously scheduled contact. (If necessary, last-minute changes in a contact plan can be accommodated in ad-hoc fashion by transmission of low-level spacecraft commands to the affected network nodes; the ION system includes functions that support this expedient.) An internationally standardized application-layer protocol for propagation of contact plan updates has yet to be defined.

Another key technical challenge to overcome in integrating DTN into mission operations infrastructures is the synchronization of clocks throughout the network. The DTN protocol implementations need to know the current time, with some accuracy, for three purposes:

1) The initiation and termination of bundle transmission and reception processing by the DTN software is driven by the scheduled times at which radio contacts will begin and end. To the extent that the current time reported by the spacecraft clock differs from that reported by clocks at ground stations or other spacecraft, the spacecraft might be transmitting when the ground station is not receiving, or vice versa. So long as this divergence is only a matter of a few seconds, automated retransmission in DTN will correct for the resulting loss of transmitted data. However, uncompensated clock drift on the order of hours could prevent altogether the utilization of scheduled communication opportunities.
2) All network administration information issued by spacecraft, such as notices of processing anomalies, significant state changes, and network activity statistics, is time-tagged to enable correlation with logged spacecraft and ground station events so that problems may be investigated and resolved. Significant errors in time tags make this analysis difficult.
3) The Bundle Protocol is designed to prevent degraded network resource utilization by automatically deleting undelivered bundles for which no further delivery efforts should be made. The basis for this determination is a user-specified "time to live". When the sum of a bundle's creation time and its time to live exceeds the current time, the bundle is deemed "expired" and may be deleted. A substantial error in the current time as reported by the local clock could result in premature bundle deletion.

Here again, there are two aspects to the problem. First, the delta between the time reported by a spacecraft clock and the correct time (as reported by a stable and highly accurate Earth-based reference clock) must be computed at fairly frequent intervals, continuously accounting for cumulative clock drift. However, for most missions this clock correlation activity is already a standard, if manually managed, operating procedure. The additional issue raised by DTN is the communication of that computed delta to the software running on the affected spacecraft, so that the software can always determine the correct current time by simply adding the delta to the time reported by the spacecraft clock. This is easy in concept, and in fact DINET employed a simple implementation of this capability. Again, though, no standardized protocol for propagation of this information has yet been defined.

VI. Operational Considerations

A factor in the schedule-enabled operation of DTN protocols is the distance (measured in light seconds) between each network node and each of its neighboring nodes – that is, all the nodes with which that node will be in direct radio communication. This figure bears particularly on the computation of the timeout intervals on which retransmission-based communication reliability depends.

Mission navigation procedures already calculate spacecraft range with high accuracy to enable the timely initiation of trajectory correction maneuvers, orbit injection, etc. A near-term means of providing one-way-light-time values to DTN software, then, would be to simply obtain them from mission navigators and transmit them in bundles. Indeed, this is exactly what was done for the DINET experiment, and it worked very well.

However, in-flight spacecraft are traveling at high speeds; while the distances to their neighbors may not change significantly (i.e., by as much as one light second) in an hour, they certainly will change at least daily. As the number of spacecraft in the Solar System Internet (SSI) grows, the chore of frequently re-computing and propagating new range values for all relevant pairs of network nodes from a central Earth-based network operations center could become time-consuming and expensive.

A more scalable approach might be to provide the spacecraft themselves with simplified navigation software and orbit computation information so that each one could individually, automatically compute the distance to each of its neighbors as required. A prototype implementation of this capability has been demonstrated[7]. Yet to be addressed are the infusion of this software into the ION system and the development of a standard protocol for propagating the orbit computation information on which it relies.

Finally, in the long term a scalable, general solution to the overall problem of cost-effective DTN network management is going to be required. The mechanisms developed for DINET might form at least part of the needed infrastructure, but operations center software that is carefully engineered for human usability must be built on that structure. The following kinds of capabilities will take on particular importance:

- A system for visualizing the highly dynamic topology of the SSI will be needed. Visualization of Internet topology is relatively straightforward, as it is mostly static. SSI topology will be fairly static in four dimensions, in that contact intervals will be almost entirely scheduled for some time to come, but in order to be usable by human network administrators it may have to be modeled by animation of a 3-D visualization.
- The DTN protocol implementations operate within a managed parametric environment that can be configured to accommodate a very wide range of mission designs. While this extensive configurability broadens the applicability of the software, it also increases the risk of error in configuration. Automated support for node configuration, with built-in sanity checking of configuration decisions, will greatly reduce the cost of SSI network administration.
- The complex interactions among DTN protocol capabilities and possible environmental anomalies make troubleshooting the network a challenge. Automated SSI fault analysis facilities, possibly leveraging JPL's expertise in machine learning and rule-based systems, could help minimize failures that compromise the performance of the network.

VII. Future Applications, Scenarios

Space internetworking via DTN offers some interesting, novel, and potentially very beneficial scenarios briefly explored in the following sections.

A. Rich Solar System Communications via Solar System "Backbone"

As noted previously in this paper, the collection of space links to sustain and support the NASA/JPL and international Mars missions tends to be manually intensive. Within the DSN, a technique known as Multiple Spacecraft Per Aperture (MSPA) has been implemented to enable multiple Mars missions to utilize the same DSN antennas simultaneously. At the current time, this means that two or more spacecraft can transmit data to Earth via one DSN antenna, but only one can have data forwarded to it via that antenna at any given time.

Although very cost effective from the perspective of serving a greater number of missions without building more antennas, this technique has inherent limitations that prevent the approach from scaling significantly. For example, the degree of coordination among various missions increases directly in proportion to the number of missions utilizing the same DSN antenna. The multi-mission coordination aspect includes, but is not limited to, swapping the single forward space link (and hence command radiation service capability) among the various spacecraft within the

[7] A brief overview of the prototype functionality is provided in the Appendix.

DSN antenna beam width and coordination of priorities to determine which mission will have the "dedicated" service should the link margin unexpectedly deteriorate for one or more of the spacecraft simultaneously being serviced, requiring that efforts be made to recover the link.

DTN-based space internetworking supports an interesting alternative approach: dedicated communications assets. The concept of dedicated communications spacecraft is well-demonstrated for use at Earth. If dedicated communications assets were to be put in orbit at Mars with DTN technology, it would be quite feasible to provide a service that is essentially continuous for those missions that make use of the communications infrastructure and is subject only to local (at Mars) communication disruptions. Rather than having a science or manned mission be concerned about continuously scheduling an Earth-to-Mars space link, a DTN implementation coupled with dedicated telecommunications spacecraft at Mars would essentially replace this set of concerns with an "always on" service subject only to physical latency considerations. Maintenance of the Earth-to-Mars space link becomes a much smaller operational concern in the presence of this "virtual trunk line." This capability would provide a complete telecommunications infrastructure from which missions could then negotiate for quality of service considerations (i.e., bandwidth and latency guarantees), as opposed to multiple link-level scheduling considerations.

B. Increased Flexibility in Mission Operations Architectures

The standardized, automated utilization of communication links that is provided by DTN network functionality enables entirely new classes of mission operational architectures. For example, DTN enables parallel end-to-end data paths to be easily exploited so that a data relaying infrastructure could be technically and organizationally heterogeneous. A spacecraft could return data via multiple orbiters – even multiple ground stations – administered by different space agencies, either alternatively or concurrently, for increased mission safety and reduced data delivery latency. In light of this, autonomous support for in-situ communication among spacecraft at a distant planet could enable event-driven, policy-based, collaborative investigation of transient phenomena without waiting for commands from Earth.

Infrastructure standardization could also enable new operational structures to be developed dynamically in response to changes in mission profiles. In place of multiple isolated and inflexible mission communications plans based on pair-wise ground tracking agreements we would have a communication fabric that can change over time, making it easy for newly arrived spacecraft to leverage the capabilities of spacecraft in extended mission.

The inclusion of DTN-enabled nodes scattered throughout the solar system would introduce a more robust communications infrastructure that is less sensitive to direct-link interruptions. Currently, for example, the operating spacecraft at Mars cease operations for about a month during the period when Mars passes near to or behind the Sun from the perspective of Earth. The solar flux from the Sun during this solar conjunction introduces such a high level of interference in the signals from Earth as they are received by the spacecraft as to introduce a high probability of corrupted or impartial command loads. Every two years when this interruption occurs, the projects lose millions of dollars idling the operations staff and the scientific missions are negatively affected, particularly those whose data collection is time-sensitive. Adding DTN-enabled nodes throughout the solar system would provide a mechanism to relay the data around the sun during these periods, thus enabling the scientific missions to continue undisturbed.

Further in the future, DTN might support mission designs that currently seem to be at the outer edge of feasibility. Investigation of caves by subterranean rovers, or investigation of a water ocean under an icy surface by an autonomous submarine, would certainly require in-situ relay communications. Deployment of a constellation of micro-satellites might enable rapid completion of a planet-wide surface investigation, but again an in-situ network providing more powerful transmission resources would be needed to link those satellites with investigators on Earth.

VIII. Conclusion

Current data relay operations at Mars rather resemble the earliest days of terrestrial computer networking when UUCP and Bitnet were prevalent; messages were passed along, hop by hop, over dial-up circuits. But even that approach, limited as it was, used agreed standards for interoperability. Our current approaches, effective as they are at returning data from Mars, are far more idiosyncratic and ad hoc. Increasing numbers of interacting spacecraft, low power missions that require data relaying, and interdependent mission scenarios are driving requirements for something that operates more like an Internet in space. This space-based Internet will need not only to provide interoperable data transfer among elements designed and operated by different agencies but also to contend with periodic connectivity, long light travel times, and a noisy operating environment.

The technical approach that is being proposed is to adopt a suite of space internetworking protocols collectively referred to as Delay Tolerant Networking (DTN). Just as the terrestrial Internet is not just TCP/IP, DTN is not just

the bundle protocol. There are other supporting elements that are part of the suite that provide functionality for routing information exchange, network management, security, and for handling peered connections between network elements operated by different agencies. This involves not only technical elements that provide the mechanisms but also programmatic elements such as service agreements, peering arrangements, integration and validation of new elements and overall governance of what is to become a new, multi-agency, infrastructure. As in the earliest days of the Internet, there are cultural and political issues that need to be addressed. These range from "that's not how we do it now" to "how can I rely on someone else's mission" to "how can I control others' use of my assets". The list of concerns goes on. Aside from reaching agreement on the right technical path, some of the biggest questions to be answered are whether these cultural barriers can be surmounted, and how.

An approach that seems a promising way to lower these real or imagined barriers to adoption is to actually demonstrate that the Solar System Internet is feasible. Experiments such as those described here, the DINET demonstration using ION on Deep Impact, ISS, and even ground-based experiments are all intended to validate the protocol's functionality, demonstrate its effectiveness, and mature the current implementations. Some of the next steps are likely to involve producing space-qualified radios that incorporate not just standard transceiver functions but also link layer capabilities, DTN networking capabilities, and sufficient on-board storage to manage relayed communications. These will need to be designed such that they can be directly integrated into a spacecraft to provide these new services. Such devices may be thought of as "network appliances", somewhat equivalent to the DSL or cable modems in many houses today that provide full access to Internet services. Of course, in order for a spacecraft, even with one of these capable radios, to provide space-based internetworking services, it will also need to provide power to the radio (possibly "always on" for demand access requests from users), to allocate communications time and radio bandwidth, and to provide any required pointing or assigned antenna time.

Once the technical specifications have been defined and implemented, and a spacecraft has been launched, there is still the small matter of coordinating and orchestrating these communications services. As at present, there is first the basic requirement to choreograph and coordinate when the involved spacecraft and ground stations will be able to see one another and when they will have resources available to perform a coordinated communications session. The existing models for doing this planning, scheduling, and coordination are rather complex and human intensive, even within a single organization. Present efforts to apply automation to both data link management and cross support planning will improve that, but further work is needed to do this in a fashion that will work across multiple mission and agency interfaces. When viewed from a DTN-based space internetworking perspective, it should also be noted that both the DSN (terrestrial to Mars space link management) and the Mars Relay Network (Mars local space link management) tend to have a significantly overlapping set of concerns, ranging from communication geometry computations to network performance extraction and reporting. In essence both these existing networks already address DTN-type management concerns, albeit not in a particularly scalable fashion.

Some new multi-agency governance model, dealing with how the infrastructure is managed and moderated, will need to be developed, attempted, and refined. Currently we have agreements between individual nodes in the network, but there is not yet a service agreement mechanism or an integrated network-wide agreement to address concerns about level-of-support, data transfer latencies, and prioritization between competing demands on network resources. This new governance model will have to define how to manage both Earth-to-spacecraft and spacecraft-to-spacecraft data links, and also how to fairly govern the transfer of data across elements belonging to multiple agencies.

However, if present successes in data relaying, improved science return, and lower power demands are any indication, the benefits to be accrued from deploying a space internetwork should be significant. Furthermore, the availability of in-space networking will permit innovative new mission configurations, and this will open up new opportunities for collaborative science and exploration, whether robotic or human.

Appendix

A. Light Time Daemon

Light Time Daemon (LTD), a recent prototype development at JPL, computes light times between two points in space-time. For a DTN network, this information is necessary information and must be available at each of the DTN nodes.

LTD's multiple operating modes support a variety of use cases. For example, LTD can be configured to operate strictly on ephemeris input data. In this high accuracy mode, the computational accuracy is limited by the input ephemeris accuracy. In the low bandwidth mode (i.e., a lesser amount of ephemeris input data), LTD can compute a fit and strictly rely on these coefficients and an interpolation function to compute the light-time to a defined level of accuracy. This is based on 13th order Lagrange polynomials that have been used successfully in operating the DSN

for several years. LTD further supports multidimensional fits, with each dimension representing a pair-wise light-time fit. By setting the accuracy-level parameter appropriately, the user can minimize the information flow needed to update remote DTN nodes. This mode of operation is appropriate for spacecraft operating as remote DTN nodes – in which limitations in bandwidth, memory and latency are primary constraints. Additionally, LTD supports piecewise updates of light-time fits. This further minimizes the bandwidth required to keep all of the DTN nodes synchronized. For example, a spacecraft in cruise mode may have a fit which spans months or even years.

The prototype is still in development, but it does appear to be promising for providing local (i.e., at the spacecraft) computations in support of DTN routing functions.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

General, Background:

1. Gladden, R., et al, "Coordinating UHF Relay Activities at Mars", *IAA International Conference on Low-Cost Planetary Missions,* Noordwijk, The Netherlands, September 2003.
2. Gladden, R., et al, "Mars Relay Coordination Lessons Learned", *IEEE Aerospace Conference #1065,* Big Sky, MT, March 2004.
3. Edwards, C. D., "Relay communications for Mars exploration", International Journal of Satellite Communications and Networking, 2007, 25:111-145.
4. Gladden, R., "Mars Reconnaissance Orbiter: The History of Supporting the Phoenix Lander", *AIAA Space 2009 Conference & Exposition*, Pasadena, CA, September 2009.
5. Edwards, C. D., et al, "Telecommunications Relay Support of the Mars Phoenix Lander Mission", *2010 IEEE Aerospace Conference*, Big Sky, MT, March 2010.
6. Gladden, R., "Mars Relay Operations Service (MaROS): Rationale and Approach", AIAA *SpaceOps 2010 Conference*, Huntsville, AL, April 2010.

Cited in this paper:

[i] IOAG web site: http://www.ioag.org
[ii] Recommendations on a Strategy for Space Internetworking, November 15, 2008: http://cwe.ccsds.org/ioag/Final%20Products/SISG%20Report%20v1.4%20FINAL.pdf
[iii] "NASA Successfully Tests First Deep Space Internet": http://www.nasa.gov/home/hqnews/2008/nov/HQ_08-298_Deep_space_internet.html
[iv] Bundle Protocol RFC: http://www.faqs.org/rfcs/rfc5050.html
[v] Licklider Transmission Protocol: http://www.faqs.org/rfcs/rfc5326.html
[vi] CCSDS File Delivery Protocol (CFDP): http://public.ccsds.org/publications/archive/727x0b4.pdf
[vii] CCSDS Asynchronous Message Protocol (AMS): http://public.ccsds.org/publications/archive/735x0b1.pdf (in formal release process as of the writing of this paper)