Evolutionary Paths to Internationally-Standardized Space Internetworking

Adrian J. Hooke [1]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, USA
Manager, Networking Architecture, Standards and Technology
Space Communications and Navigation Office, NASA Headquarters, Washington DC, USA

The space community is now poised at the start of an exciting new era of space exploration, where many different organization will be launching spacecraft to investigate nearby bodies, such as the Moon and Mars. International fleets of spacecraft are thus likely to begin to fly and interoperability will often become a mission-enabling technology. These fleets will require a mix of powerful data communications capabilities to be available, ranging from conventional point-to-point architectures to new configurations supporting internetworked operations. This paper reviews the various CCSDS capabilities to support fully-standardized in-space internetworking, and outlines a strategy for smoothly evolving into the new era with minimum impact on the highly successful installed base of CCSDS-compatible systems, which are now widely deployed across the international space community.

I. Introduction – The Early Days of Space Data Communications

The space age is now roughly fifty years old and has been approximately bisected by the era of standardization of space mission operations, which began in 1982 with the formation of the Consultative Committee for Space Data Systems (CCSDS). Up to that time virtually all missions flew customized flight and ground data system designs, with their Telecommand (TC) and Telemetry (TM) systems being simple digital versions of the analog systems flow on early missile systems. For TC, the digital data structures were primarily discrete (and asynchronous) commands that were transmitted as unique bit patterns within a synchronous forward link symbol stream. For TM they were complex Time Division Multiplexed (TDM) formats of Pulse Code Modulated (PCM) data consisting of a main commutation frame of measurements (Figure 1) that were sampled in fixed rotation. Often, one or more of those sampling slots would have a sub-commutator running, acquiring yet another cycle of measurements sampled at a low rate.

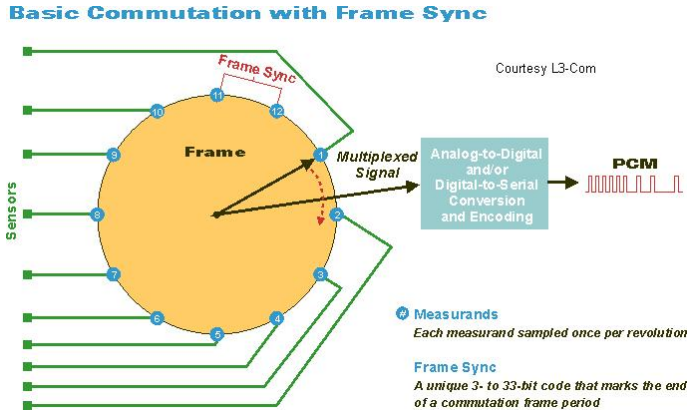

Figure 1: Early Commutated TDM Telemetry

These TDM streams emulated the motor-driven commutators of the early missiles, where analog measurements were mechanically sampled and directly modulated the return link. Inherently, they shackled every source of data to a fixed time slot for transmission and thus intimately coupled the application with the synchronous communications link. Often, they resulted in significant oversampling of the source and thus wasted data transmission capacity.

[1] ☐ This work was performed in part at and under the direction of the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration. Copyright © 2008 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U.S. Government has a royalty-free license to exercise all rights under copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

In the mid 1970's, this tightly coupled data communications model was broken by the introduction of packetized data transfer, where sources of data were allowed to create a autonomous, variable length "Packets" of information. First proposed by Greenberg and Hooke at JPL in 1974 as a New Technology Report[1] for "Block Telemetry" and independently by Greene at GSFC in 1977 as an "Asynchronously Multiplexed Telemetry System"[2], these converged by 1978 into a draft NASA "Space Data Packetization Standard"[3]. Shortly thereafter a "NASA-ESA Working Group" was formed to begin the internationalization of the new standard and in 1982 the CCSDS was created to complete the initial standardization of "Packet Telemetry" and "Packet Telecommand" (Figure 2).

These internationally-agreed CCSDS Packet standards spelled the end of the 25-year reign of TDM and heralded a new era where packet switching of spacecraft traffic became the predominant technology for the implementation of new civil space missions.

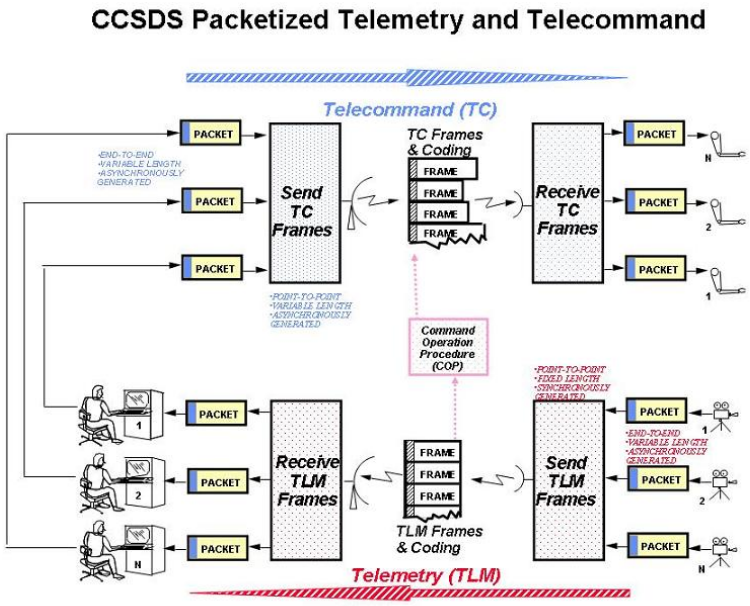

Figure 2: CCSDS Packet TM and TC, circa 1982

In 1974, a seminal paper[4] by Cerf and Kahn was published that outlined the basic concepts of TCP/IP (which later grew into the foundation for the Internet). This work was based on the early concepts of Leonard Kleinrock[5], who developed the basic theory of packet switching as part of his Ph.D. thesis. So in 1974, the early concepts of packetized data transfer were independently forming in both the terrestrial and space data communications communities. It would be almost exactly twenty five years later that these concepts again converged into the "Interplanetary Internet".

II. The Beginnings of Space Networking

By the mid 1980's, CCSDS was vigorously pursuing the additional international standardization of new space/ground data communications techniques, including Modulation, Coding and Data Compression -- itself enabled by the standardization of the new high-performance Reed-Solomon coding. When plans emerged circa 1985 to build an international space station, CCSDS formed two special working groups to address its unique challenges:

a. New requirements for the standardization of a space/ground data link interface characterized by high forward data rates, very high return data rates and the need to handle many different types of traffic – including audio and video to support flight crew activities.
b. A new need to support internetworked operations for the transfer of user information between the spacecraft and the ground.

The emergence of "internetworking" in the early days of the space station was an interesting development, because the design of the mission operations system was centered on a space-based local area network (using the Fiber Distributed Data interface, FDDI) that was paired with an open internetworking ground environment. Both were to be interconnected by the high-capacity space/ground data links. At that point in time, the concept of "Open Systems Interconnection" (OSI) was emerging and the International Standards Organization (ISO) was heavily engaged with the development of international standards in accordance with the 7-layer OSI model. Four standards – part of the US "Government Open Systems Interconnection Profile" (GOSIP) were of particular interest:

- The Connectionless Network Layer Protocol (ISO 8473)
- Its accompanying Network Layer Security Protocol
- The OSI Transport Protocol (TP4)
- The File Transfer Access and Management protocol (FTAM)

For over two years, CCSDS and the space station program worked intensively on what became the CCSDS Recommendation for "Advanced Orbiting Systems" (AOS). The result was the infamous "Pipe Diagram" of circa 1990 (Fig. 3), which established a rich set of point-to-point services across the space link (including the ability to segregate different data flows into different 'Virtual Channels") and which also recognized two forms of end-to-end internetworking:

- CCSDS Path Service, based on the conventional CCSDS Packet of Packet TM/TC fame, which exploited relative static routing relationships by defining a connection-oriented protocol for user data transfer.
- A connectionless "Internet" service, rooted in ISO- 8473 at the Network layer.

For transfer of IOS 8473 datagrams across the space link, the AOS required the definition of a new Encapsulation Service which simply wrapped the 8473 packets inside CCSDS Packets and thus allowed both CCSDS Path and Internet services to share the same physical channel. This important principle of insulating the space links from the varied and changing technologies of the upper layers, by using an Encapsulation mechanism, was thus established. In subsequent years, the Encapsulation Service was broken loose from its dependence on the CCSDS Packet and a new, streamlined Encapsulation Packet emerged. But it would be almost twenty more years after the international standardization of AOS and its adoption by the Space Station program before Encapsulation returned to prominence.

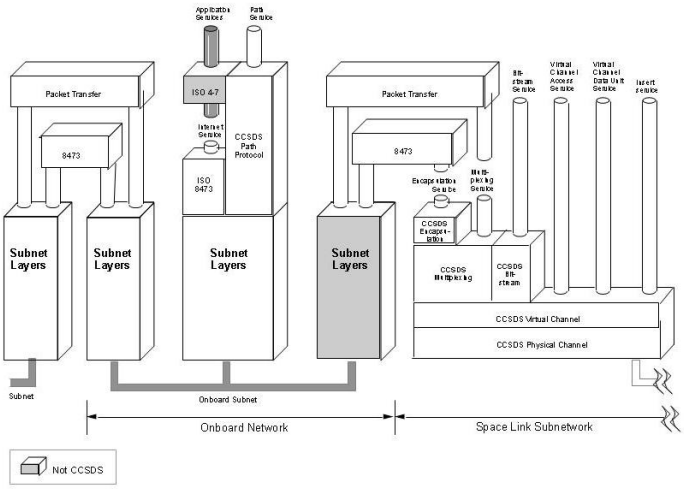

Figure 3: The AOS "Pipe" Architecture, circa 1990

By the mid 1990's, the Space Station was in the middle of an extensive series of redesigns but the AOS standard remained the bedrock of international interoperability. Meanwhile, a clash of the titans was in progress between the GOSIP protocol suite and the emerging Internet suite - based in IP and including TCP/UDP and FTP - as the replacements for ISO 8473, TP4 and FTAM. During this period of turmoil, new space mission concepts began to emerge that featured large constellations of spacecraft requiring highly efficient in-space data networking. Prominent among these were the "Brilliant Eyes" and "Brilliant Pebbles" of the US Strategic Defense Initiative (SDI).

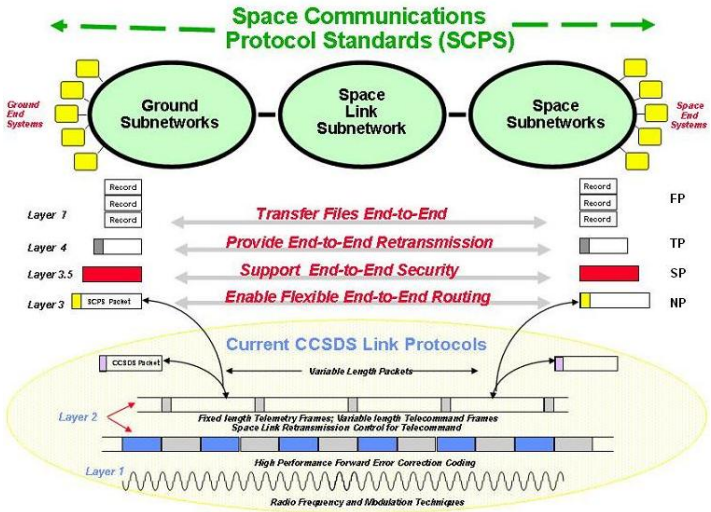

Figure 4: "SCPS", circa 1996

Recognizing the potential importance of the emerging Internet technologies to such missions, NASA and DOD teamed up in 1994 to extend the IP suite into space. The result was the "Space Communications Protocol Standards" (SCPS) project, which began the development of new standards that could enable terrestrial Internet to be extended for use by the space mission operations community. This work soon transferred into CCSDS (Figure 4) with the goal of producing international standards.

The SCPS development focused around four layers of protocol:

- A SCPS Network Protocol (NP), using IP addressing but implementing very low-overhead packet headers in order to conserve space channel communications resources.
- A SCPS Security Protocol (SP), implementing basic data protection services but again using low-overhead headers.
- A SCPS Transport Protocol (TP), which extended the Internet's TCP to allow it to work effectively over noisy, long-delay space links.
- A SCPS File Protocol (FP), which extended the Internet's FTP to allow efficient space/ground file transfer over long-delay, disrupted space links.

The FP and the TP components of the SCPS suite were designed for full interoperability with their native terrestrial Internet counterparts. The SP and the NP were not intended to be directly interoperable with IPSec and IP, requiring some sort of translation gateway in order to complete an end-to-end IP-like dialog. As part of the development of the SCPS suite, the protocols were flown in space on the U.K. "Space Technology Research Vehicle"[6] in early 1996. During these tests, STRV 1b was given its own IP address – possibly the first spacecraft to do so, and also the first to demonstrate that telemetry and telecommand data could be sent using Internet technologies from remote ground sites to an operational spacecraft.

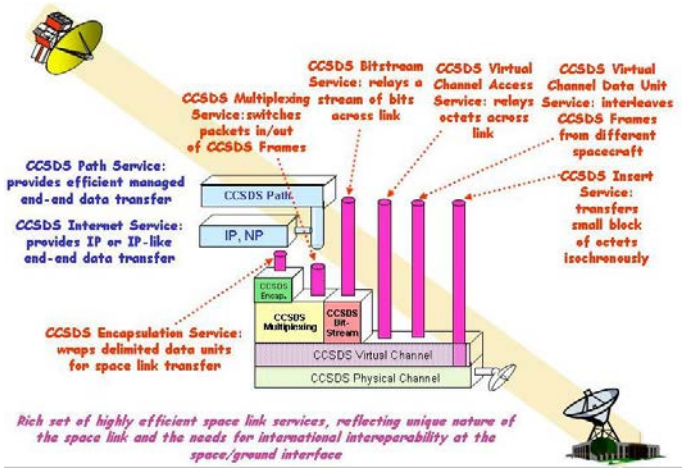

Figure 5: The "AOS Pipe", circa 1998

By 1998 the AOS "Pipe" diagram was being re-drawn (Figure 5) to show IP and NP as different types of Network layer "Internet Service") data transfer over the CCSDS space link, in parallel with the widely-deployed CCSDS Path Service and as an end-to-end alternative to the point-to-point Link layer services.

III. The Interplanetary Internet

In 1998, one of the "fathers of the Internet" – Dr. Vint Cerf – joined the space community as a JPL Distinguished Visiting Scientist, with a principalgoal of defining how to expand the terrestrial Internet across the Solar System.

The potential for achieving convergence between the problems of space and terrestrial communications – the confluence of "two rivers" of technology (Figure 6) - was recognized by the US Defense Advanced Research Projects Agency (DARPA), who allocated funding from their Next Generation Internet program to study the possible architecture of an InterPlaNetary Internet (IPN).

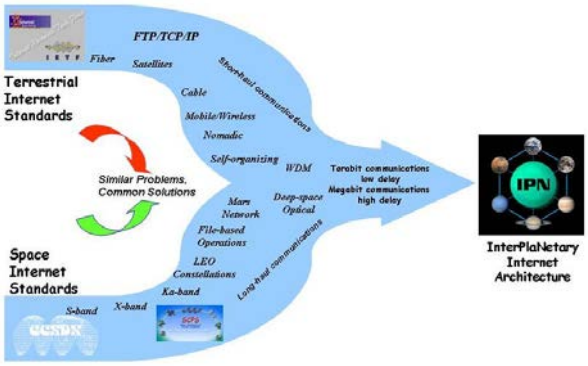

Figure 6: "Two Rivers" leading to the Interplanetary Internet

The end goal was the kind of interplanetary communications scenario shown in Figure 7, where many local "islands of connectivity", dispersed across the Solar System, could be interconnected by a backbone of long-haul space links. Under the sponsorship of DARPA, a new "Delay Tolerant Networking (DTN) Research Group" (DTNRG) was formed in the Internet Research Task Force (IRTF) to define the necessary architecture and protocol suite.

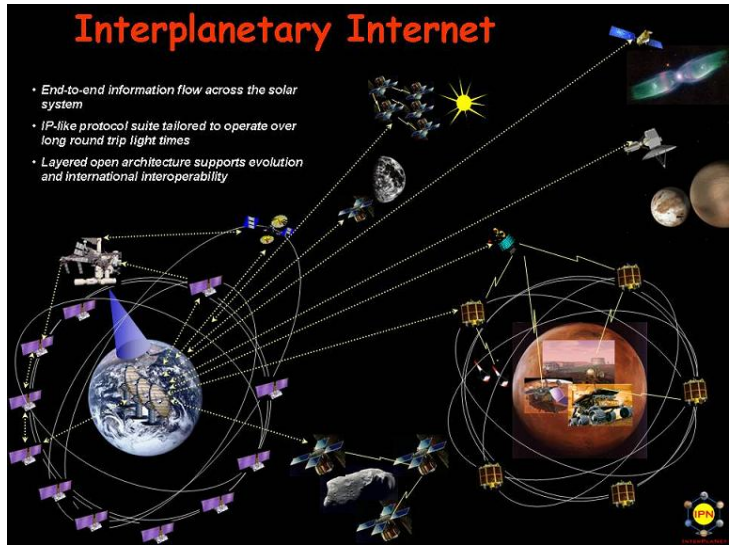

Figure 7: The Interplanetary Internet

The DTN architecture[7] that emerged was that of an overlay network based on a new "Bundle Protocol" (BP), where the relationship of BP to IP is as IP is to Ethernet, i.e., a TCP connection (perhaps the entire Internet) within an IP-based network may be one "link" of a DTN end-to-end data path; a deep-space radio transmission may be another. The BP itself (Figure 8) therefore sits below the end-to-end space applications and above sub-networks that may be the terrestrial Internet, the individual long-haul CCSDS backbone space links, or local area networks at remote locations in space.

"Delay Tolerant Networking" has been generalized into "Disruption Tolerant Networking" (network disruption is, in essence, highly variable delay). Two DTN-friendly applications were already in work with CCSDS – the first (in 1999) being the CCSDS File Delivery Protocol (CFDP) and more recently the Asynchronous Messaging Service (AMS).

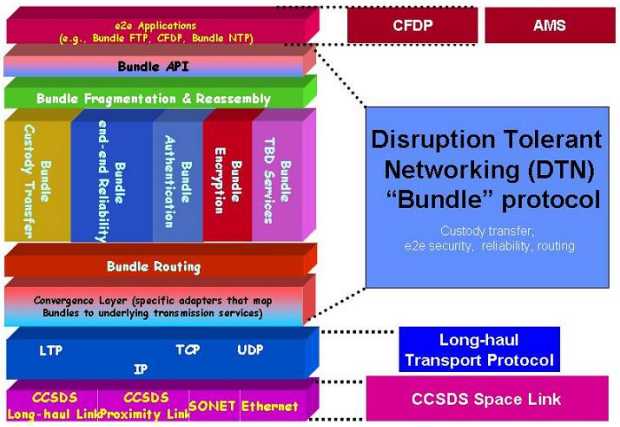

Figure 8: The DTN Bundle Protocol

The BP provides a common platform of Interplanetary Internet services to these end-to-end applications, including reliability, security and routing. Two key concepts are inherent to DTN operations:

- Store-and forward
- Custodial transfer

Within a DTN data flow, Bundles of application data work themselves through the network as a series of individual hops. At each hop, a Bundle agent takes "custody" of an incoming Bundle and assumes the responsibility to transfer it to another Bundle agent across the next hop. If the next hop isn't immediately available, the Bundle agents store their custodial cache of Bundles until the next link becomes available. In order to avoid undesirable end-to-end retransmissions, a "Long-haul Transport Protocol" (LTP) can sit above each of the individual hops to provide local reliability. LTP is also known as the "Licklider Transport Protocol", in honor of J.C.R. Licklider[8], an early data communications pioneer who's "Intergalactic Computer Network" has often been regarded as the first reference of the capability that would eventually become the Internet.

IV. The Upcoming Twenty Five Years: Building the Interplanetary Network

The space community now stands poised at the start of a new era where in-spaced networking will assume increasing importance for efficient mission operations. Already, an early international "Mars Network" has been constructed by adding data communications relay payloads – using the CCSDS Proximity-1 protocol – to scientific

orbiters operated by both ESA and NASA. Approximately 95% of the data returned from landed Mars vehicles is now transferred via these orbiting relays. With an increased international interest on missions to the Moon, scenarios have been constructed (Figure 9) that have identified the problems of communicating with and among fleets of spacecraft operating in cislunar space.

Clearly, as the number of vehicles and their end users increases and the configuration of the data links which interconnect them becomes more complex, it soon becomes a daunting task to configure an end-to-end data flow by manually setting-up all of the individual data paths. As with the Earth's Internet, the solution is to begin using the power of a Networking protocol.

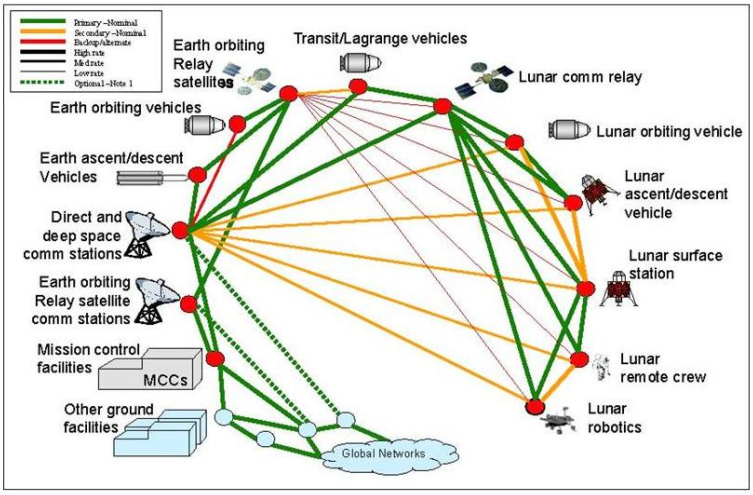

Figure 9: Cislunar Networking Scenarios

The current CCSDS protocol suite (Figure 10) is well positioned to support this smooth evolution by building upon the current robust and extensively-deployed underpinning of CCSDS Physical and Link layer protocols. These include a wide selection of internationally standardized modulation techniques, advanced channel coding schemes - including new emerging standards for very high performance Low Density Parity Check (LDPC) codes – and the four "workhorse" space Link protocols: Telemetry; Telecommand; AOS; and Proximity-1. Ground access to space links using these protocols is supported by the CCSDS "Space Link Extension" (SLE) services.

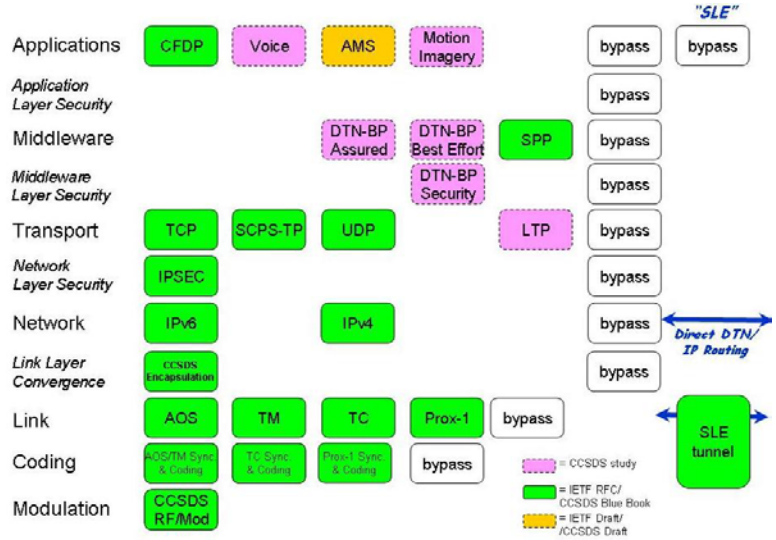

Figure 10: The Emerging CCSDS Protocol Stack for Space Networking

Sitting above the CCSDS space Link layer protocols are a variety of upper layer capabilities. The CCSDS Space Packet Protocol (SPP) continues to be the basis for user data transfer in a large number of current and future missions, with the CCSDS File Delivery Protocol (CFDP) now emerging as a protocol of choice for custodial, store and forward transfer of files to and from the spacecraft. A new Asynchronous Message Service (AMS) standard is under development, to support discrete message transfer as an alternative to file transfer. Development of DTN and LTP is well underway within the IRTF and its profiling for space use is just starting within CCSDS. New CCSDS standards for voice and motion imagery are under study.

Use of the Internet protocol suite – based in IPv4 or IPv6 - is possible as a result of the SCPS development and the refinement of the CCSDS Encapsulation Service to provide a convergence layer between IP and the CCSDS space links.

The CCSDS suite is now well postured and configured to allow three basic options for providing Network layer service to space mission users (Figure 11):

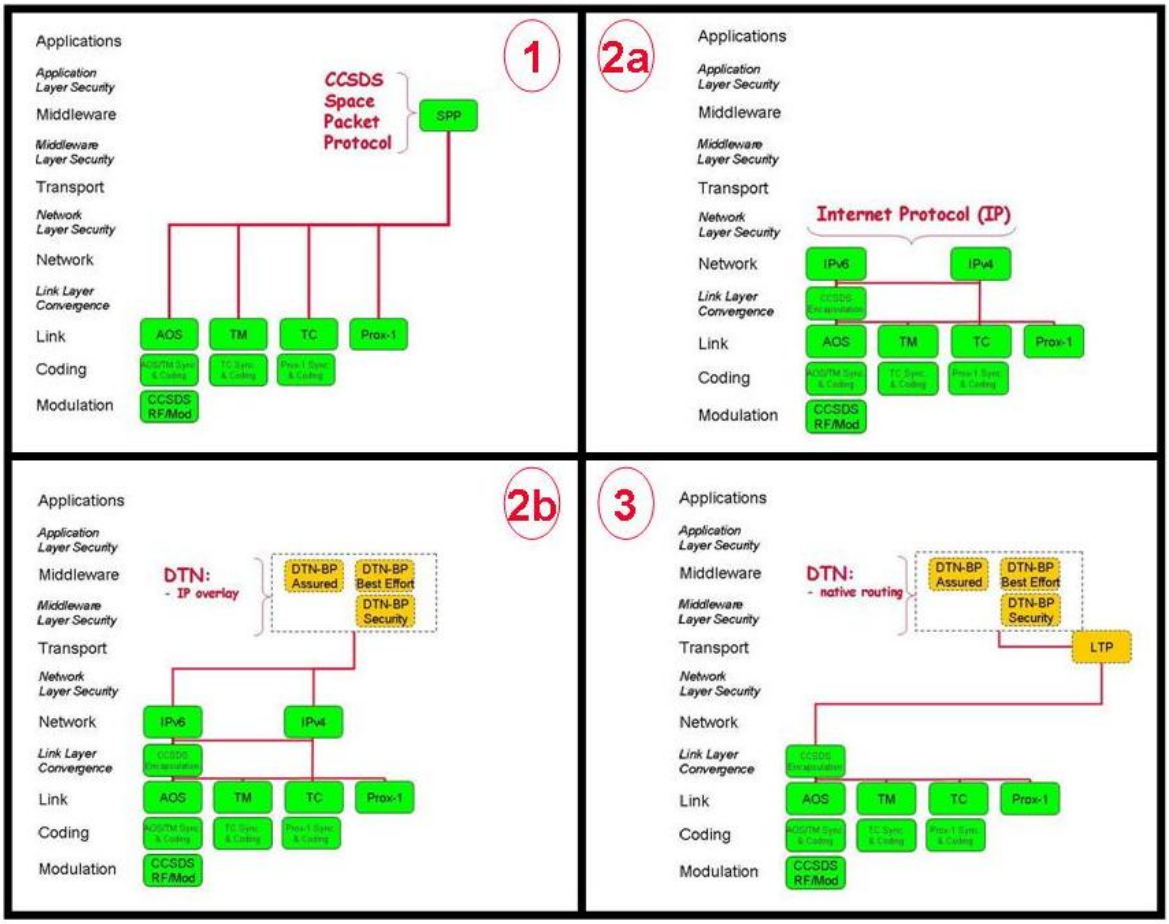

Figure 11: CCSDS Options for Space Networking

- Option 1: the conventional CCSDS Space Packet Protocol.
- Option 2: "IP" routed service, using either version 4 or version 6 of the Internet Protocol.
- Option 3: DTN routed service, using the Bundle Protocol augmented as necessary with LTP.

Option 1, which has been widely deployed with great success, is expected to continue to provide valuable service to missions with relatively simple configurations and low demands for dynamic routing.

Option 2 - native IP routed space networking- is primarily envisaged to be selected in relatively short delay, richly connected, "chatty" data communications environments, as would be expected in cislunar space. The "2b" sub-option allows DTN to run as an overlay network on top of IP routed service, providing capability to recover from network partitioning caused by disconnections.

Option 3 – native DTN routed space networking – is expected to emerge as the preference for long-haul interplanetary internetworking where it would be used to interconnect local islands of Internet or Internet-like connectivity, such as on remote spacecraft or on the surfaces of other Solar System bodies.

V. Conclusion

The space mission community is now well equipped to take the next big leap towards fully internetworked space mission operations. A robust suite of internationally standardized CCSDS protocols forms the springboard for this evolution, providing a clear way forward in a way which is completely compatible with the large current installed

base of CCSDS systems and which is fully backwards-compatible with the investments that have already been made.

References

[1] New Technology Data Sheet #3391, Jet Propulsion Laboratory, Edward Greenberg and Adrian J. Hooke "A Spacecraft Block Telemetry System Compatible with Ground Transmission Circuits", 05 December 1974.

[2] Edward P. Greene, "A Proposed Asynchronously Multiplexed Telemetry System for Spaceflight Operations", Goddard Space Flight Center, August 1977

[3] Aerospace Data Systems Standard 3.3, "Space Data Packetization Standard", Goddard Space Flight Center, 31 May 1978

[4] Vinton Cert, Robert Kahn, "A Protocol for Packet Network Intercommunication", IEEE Transactions on Communications, May 1974.

[5] Leonard Kleinrock, "Information Flow in Large Communication Nets", Ph.D. Thesis Proposal, Massachusetts Institute of Technology, July 1961

[6] Cdr Richard Blott RN and Nigel Wells, "THE SPACE TECHNOLOGY RESEARCH VEHICLES: STRV-1A, B", Small Sat conference, 1996.

[7] Request for Comments (RFC) 4838, "Delay-Tolerant Networking Architecture", V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, H. Weiss: April 2007

[8] Licklider, J. C. R. "Topics for Discussion at the Forthcoming Meeting. Memorandum for: Members and Affiliates of the Intergalactic Computer Network", Advanced Research Projects Agency, 23 April 1963.