76th International Astronautical Congress (IAC 2025), Sydney, Australia, 29 Sep-3 Oct 2025.

IAC-25,B2,5,9,x101394

Internet Protocol Stack in Deep Space: Architecture and Simulation Results

Marc Blanchet
Viagénie
Québec, QC, Canada
marc.blanchet@viagenie.ca

Wesley M. Eddy
Aalyria Technologies
Avon, OH, USA
wes@aalyria.com

***Abstract*—This paper describes the use of the Internet Protocol stack to implement a network in deep space. The architecture is presented with a list of key needed adaptations: (1) temporary buffering of IP packets on forwarders with intermittent links, (2) configuration of the QUIC transport protocol with parameters related to the expected round-trip time and (3) configuration of application timeouts. The simulations described in this paper confirmed the suitability of the IP suite in deep space. Applications and network services are also discussed. Finally, benefits of using the Internet Protocol suite are described.**

***Index Terms*—deep space, Internet Protocol, IP, networking, delay-tolerant networking, DTN, QUIC**

I. Introduction

Deep space has been mostly limited to government agencies such as NASA, ROSCOSMOS, ESA, JAXA, CNES, CSA and others, primarily due to the very high costs and risks associated with space missions. This is starting to change, with multiple commercial lunar missions and plans for commercial asteroid and Mars missions. Government mission data systems have communicated point-to-point between terrestrial infrastructure and spacecraft. Relay has been done at the link layer or physical layer and does not use routed networking technology. For example, Fig. 1, illustrates the use of low-layer relays around Mars that have been very successful in supporting relaying for the low number of Mars surface missions. At a small scale, this has sufficed, but is not able to meet the needs of present missions in design and planning. For this reason, routed networking is now needed to replace the point-to-point paradigm and to make more efficient use of infrastructure, communications opportunities, power, and spectrum.

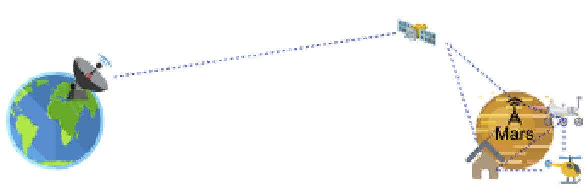

Fig. 1. Mars relay and surface network (not at scale)

The round-trip time (RTT) from Earth to Mars ranges from a few minutes to 44 minutes, due to the propagation delays of radio waves in space and the position of the planets relative to each other. Planned communication interruptions can range from a few seconds to several days or weeks, due to satellite orbital motion, planetary rotation blocking lander/rover visibility, and occlusions or conjunctions with other bodies. The orbiters acting as relays around Mars such as Mars Reconnaissance Orbiter (MRO), Odyssey (ODY), Mars Atmosphere and Volatile Evolution (MVN, Maven), and Trace Gas Orbiter (TGO), have storage to temporarily store bits during interruptions. Communications are planned using the MAROS broker software [15], as links are dedicated to a mission during a specific period of time. If a planned communication window cannot be used for one reason or another by a mission, that window is lost. If bandwidth is not fully used by a mission, then the unused bandwidth is lost. The worst case for RTT for Mars is when Mars is on the other side of the Sun from Earth, as illustrated in Fig. 2, creating a two-week period where direct communication is not possible, a situation that occurs every two years [17].

Fig. 2. Mars solar conjunction (not at scale)

A. Terminology

In this paper, we use the term ”deep space” to clearly differentiate with typical Earth orbiting missions (e.g. LEO, MEO, or GEO), where routed networking protocols and the use of IP is already well established. In this paper, ”deep space” means communication between Earth and missions orbiting or landed on the Moon, Mars, or other planets, asteroids, or comets, as well as spacecraft in interplanetary trajectories or at Lagrange points.

B. Mars Round-Trip and Holding Times

A study [16] [28] of the potential communication windows and actual events that occurred between the Mars orbiters and the rovers, and between the orbiters and the antennas on Earth, was conducted based on data provided by the JPL MAROS project [15]. In summary, the average complete path time (Earth-to-orbiter-to-rover-to-orbiter-to-Earth) is 12 hours and 40 minutes, the maximum is 161 hours, and the minimum is 79 minutes. The average holding time of data in orbiters is 10 hours and 6 minutes, the maximum is 152 hours, and the minimum is 7 minutes.

TABLE I
MARS ROUND-TRIP AND HOLDING TIMES

	Complete Path Time (RTT)	**Data Holding Duration**
Average	12h 40 min	10h 6 min
Minimum	79 min	7 min
Maximum	161h 42 min	152 h

Although the Moon is only a few seconds away from Earth, intermittence of communications through orbiters remains a concern. It is possible for a small constellation [13] of satellites / orbiters to eliminate intermittence by providing total coverage to the lunar regimes, similar to how satellite coverage blankets the Earth. However, this has not been achieved yet in reality and will be costly and complex, so intermittent support is a network protocol stack design requirement.

II. MOON AND MARS COMMUNICATIONS AND NETWORKING ARCHITECTURE

Typically, space missions use CCSDS link layer protocols (e.g. TC, TM, AOS, USLP, or Prox-1). In the future, as shown in Figure 3, on the surface or in orbits around celestial bodies, terrestrial wireless technologies such as 5G / 6G and IEEE 802 WiFi are expected to be used [12] [11] [19] [9], in addition to Prox-1 and AOS or USLP. Deep-space links are still expected to use AOS or USLP CCSDS link layer protocols in the future. DVB-S2 and successors may also have a role in commercial architectures, due to the availability of hardware/software and present use in Earth satellite services. All CCSDS link layer protocols support the IP packet payload.

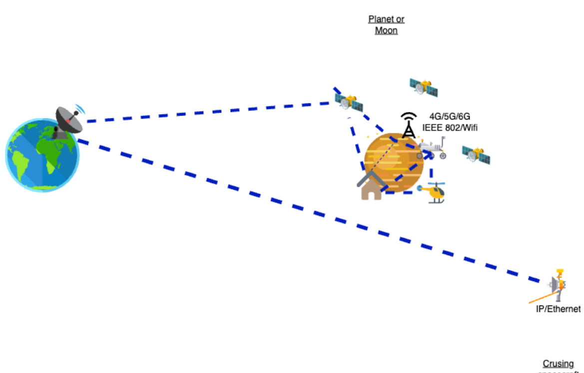

Fig. 3. Celestial body and spacecraft communications architecture

Figure 4 shows the resulting IP network. Using IP, any node on the network can be reached from any other node.

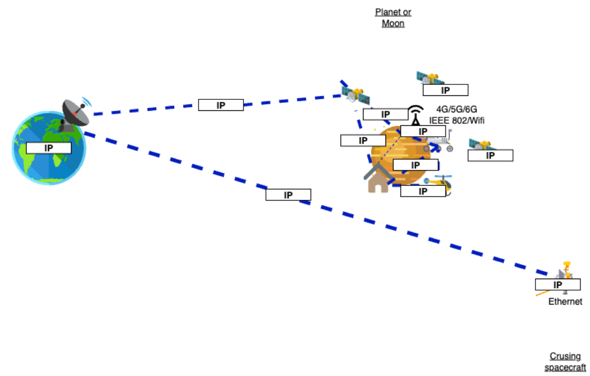

Fig. 4. Celestial body and spacecraft networking architecture

An IP-based layer 3 network enables the full use of bandwidth on all links, with dynamic allocation to different applications, faster end-to-end data delivery, and providing higher availability with alternate paths. Used in combination with a reliable layer 4 transport such as QUIC [24], it frees up the application about the knowledge of link characteristics and network or link issues such as packet loss, reordering, and duplicates.

III. RELATED WORK: BUNDLE PROTOCOL

In the early 2000s, a report [23] concluded that the IP protocol stack was not suitable for deep space networking given that its primary transport protocol, the Transmission Control Protocol (TCP), and applications were too chatty, require multiple round-trip times to establish connection, and were not geared toward long delays and disruptions.

As a result, over the last 20 years, a completely new networking protocol stack, based on the new Bundle Protocol (BP) [27], also known as delay-tolerant networking (DTN), has been designed. BP operates as an overlay or underlay with IP networks, but is not compatible with IP. Some of the main characteristics of BP that distinguish it from IP are as follows:

- The unit of data, a bundle, is stored when the next hop is not reachable, and forwarded when reachability is restored
- Node identifiers are permanent, as they are not related to topology or dynamically allocated like IP addresses.
- End-to-end reliability is not implemented.

Despite 20 years of engineering work on BP, wide commercial relevance has not been achieved, especially compared to IP. BP continues to have shortcomings that adopters need to address [5]. First, BP does not implement the equivalent of IP transport semantics, which means that a BP stack does not handle end-to-end loss, duplication, reordering, and flow control, all necessary to provide end-to-end reliability. Without these, developing applications over BP requires every application to essentially reinvent these common transport capabilities themselves. In contrast, IP-based applications that use reliable transport, such as TCP or QUIC over IP, can rely on the transport layer to ensure reliable delivery of application data. Additionally, no significant applications have emerged on

BP. Typical demonstrations use BP reachability tests (”ping”) and raw file transfer, while on the IP stack there are a multitude of classes of mature applications (web, email, voice/video, messaging, pub/sub, file sharing, etc.). Another challenge of BP is scaling with mobile nodes. With permanent BP node identifiers, a mobile BP node keeps its identifier while attached to a new network, requiring every other node on the network to be updated to the new location of the node, which does not scale and is not reliable given the delays and disruptions. In contrast, the QUIC transport protocol supports node mobility.

Other ecosystem challenges remain for BP. For example, there is a large IT industry built around IP networking that provides economies of scale. Competent IP network engineers and technicians can be found in any area, while in contrast, staffing with BP/DTN expertise is very hard to find. IP forwarding and other processing (e.g. firewalls, QoS, etc.) are widely implemented on efficient and affordable hardware. BP is generally implemented as a software system, and high-performance demonstrations use expensive platforms. Security for IP networks is well understood, and there are well-established tools such as HAIPE devices, cross-domain guards, etc. that can support high-security system needs. Similar security capabilities for BP networks are not yet available.

IV. Reassessment of IP for Deep Space

Since the initial assessment in the early 2000s which concluded that IP was not suitable and BP needed to be designed, the IP stack has evolved rapidly. Key aspects include:

- IP is now the basis of cellular systems and is used by mobile phone applications which were only emerging during the earlier assessment. IP has scaled down further than this even, now also serving as a basis for the Internet of Things (IoT) which has many similarities with deep space needs, such as low bandwidth, low energy, intermittent communications, low CPU resources, and low memory.
- IP routing technology has continued to evolve, with Software-Defined Networking (SDN) and Segment Routing (SR) as significant developments, and MPLS has become a dominant technology in efficient high-performance systems, including LEO satellite networks. These provide new tools that can be applied today for operating very complex IP-based deep-space networking systems.
- Since the early 2000s, the QUIC transport protocol [24] has been invented and accounts for a significant part of Internet traffic [1]. QUIC has a number of features that are well suited to deep space adaptation and is discussed in more detail in the next section of this document.
- IP support is already present in the platforms / operating systems being used in deep space systems. These systems have become small networks or LAN environments themselves, and typical missions now have multiple onboard processors that need to communicate within a single vehicle or rover. Extending or interconnecting these pockets of IP networking may be a more natural approach than building overlays with entirely different protocols, such as BP.

In this new context, a reassessment [4] of the use of the IP stack in deep space has concluded that the IP protocol stack is indeed suitable for deep space networking. Industry interest has been building further in this direction, and it was the subject of several IETF side meetings and a bird-of-a-feather session, resulting in the formation of a working group for Taking IP To Other Planets (TIPTOP) [29].

V. Architecture: Considerations for implementing IP protocols in space environments

The primary challenges for networking in deep space are long delays and intermittence/disruptions. There are many other challenges for communications in space, such as weak signals, error detection and correction, efficient modulation, etc. But these are managed at their respective levels in the physical and link-layer technologies.

The key adaptations needed for the IP protocol stack (network layer and above) in deep space are at the IP, transport, and application layers [8].

A. IP Layer

At the IP layer, an IP packet does not have any notion of time; therefore, an IP packet can live for a long period of time. However, a normal IP forwarder/router that receives a packet in which the destination address is not reachable based on its forwarding table simply drops the packet. In deep space, orbiters forwarding IP packets and space edge routers will face intermittent communications, which means that they should buffer the packets instead of dropping them, similar to the BP store and forward design. Policies can be established for buffering or forwarding based on various criteria, such as source or destination addresses or prefixes, Diffserv/traffic class field, flow label, or any other packet aspects. Tools for implementing those policies are already available on all common platforms.

B. Transport Layer

At the transport layer, above IP, the User Datagram Protocol (UDP) transport has no notion of time, therefore an application using UDP over IP in deep space can carry its data over a very long time period. However, UDP itself does not provide any reliability service.

The TCP transport, which provides end-to-end reliability, is not suitable for deep space, as found in the report cited in Section III. This fact remains true and the perspective to update is that, while TCP was the dominant Internet transport in the past, there are other options in the Internet stack now such as QUIC, and TCP is becoming less popular for new applications.

The more recent QUIC transport, which also provides end-to-end reliability like TCP, has many properties suitable for deep space. It is modular, configurable, and efficient, requires only 1 round-trip time (RTT) to establish the connection and the security context, supports mobility, is secure by default, can support multiple applications over a single connection, has no head of line blocking issues, etc.

However, since its primary usage is for the Internet, where connectivity is highly available and delays are very low, the default configuration of the QUIC stack is not suited for space. QUIC can be used in deep space with the proper configuration at the establishment of the connection [10]. A QUIC connection can live for a very long time. For example, QUIC connection context can be established while the spacecraft is on the launching pad and remain in use for the whole mission.

As a reliable transport, QUIC automatically manages issues on links and in networks such as packet reordering, duplication, loss, and pacing, freeing the application to take care of these issues.

There are challenges related to other transport mechanisms, such as congestion control algorithms, path MTU discovery, etc., that rely on timely feedback to measure different aspects of the path and adjust sending behavior to the current path properties. In this paper, these are assumed to be mainly dealt with suitable configuration via management, but different algorithms and approaches are viable future research and subject for discussion in standards groups.

C. Application and Application Protocols

Typical design of Internet applications assumes that the connectivity will be pretty available, bandwidth is high, and latency is low, especially compared to deep space characteristics. Therefore, most Internet applications are unlikely to work as is in deep space. However, given the use of reliable underlying transport protocols, most applications or applications protocols typically only implement timeouts at the application layer. Therefore, by configuring the timeouts with values related to the expected RTT of a mission and using an asynchronous request-response pattern, typical IP applications become suitable for deep space.

The most widely used application protocol nowadays is HTTP [25], with the latest version HTTP version 3 (HTTP/3 or H3) running over QUIC [26]. HTTP has no notion of time. There are very few HTTP headers, mostly related to caching, that contain time. These should either not be used in deep space or the value should be set appropriately. Prior HTTP versions that run over TCP may be viable in limited cases for deep space (e.g. well-connected lunar use), but are not generally suited or easily adapted to other deep space cases. Therefore, HTTP over QUIC / H3 is a solid foundation for deep space applications.

Application and application protocols that use UDP transport, such as streaming, are also suitable for deep space, given the proper configuration of relevant application settings, such as timeout.

VI. Simulation Results

As the reassessment shows, the IP stack is suitable for deep space if properly adapted. The following sections describe adaptations for the following functionalities: forwarding, transport, and applications. Simulations are also presented.

A. Forwarding and Buffering IP Packets

As a proof of concept, buffering of IP packets is implemented with 200 lines of C code, in a prototype using the Linux "tun" interface [14] and a simulation shows [3] that it is able to store packets during a link interruption and that the packets are then pulled from storage and forwarded when the link is reconnected, using a simple FIFO policy. ICMP ping is used for testing. The network consists of three nodes: a client node, a server node, and a forwarder node, as shown in Fig. 5.

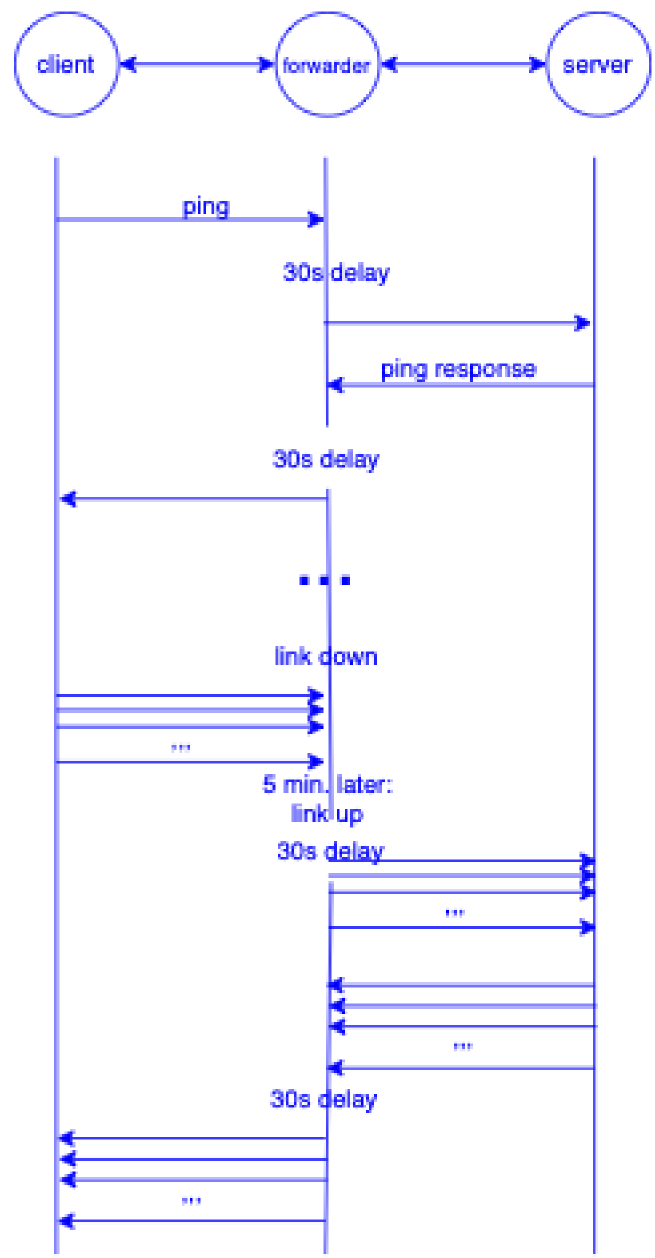

Fig. 5. Simulation of IP packet buffering using ping and 3 nodes

A 30 second delay is set on the forwarder node in both directions. At the beginning of the simulation, only the induced delay is involved; therefore, the ping client receives the replies after 60 seconds. The forwarder node then has an interruption of its link to the server node for a period of 5 minutes. The ping client does not receive any response or error during that period of time. In a normal Internet forwarder environment, the forwarder would have dropped the packet and maybe sent an ICMP destination unreachable error message to the ping client. However, in this implementation for deep space simulation, the forwarder stores the packets. After the 5-minute period, the link is back and the forwarding resumes. As shown in Fig. 6, the ping client starts receiving replies with an RTT of 360 seconds, which is 300 seconds for the link interruption and 60 seconds for the induced delay. The next replies were 359 seconds, 358 seconds RTT down to 60 seconds. This is normal as, by implementing a FIFO policy, the oldest packets were the first ones to be forwarded. Note that while the times used in the simulation are small compared to deep space times, the same results were confirmed by larger times, such as 1 hour, in simulations.

- PING fc00:1::3 (fc00:1::3) 56 data bytes
- 64 bytes from fc00:1::3: icmp_seq=1 ident=12259 ttl=62 time=60015 ms
- 64 bytes from fc00:1::3: icmp_seq=2 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=3 ident=12259 ttl=62 time=60004 ms
- ...
- 64 bytes from fc00:1::3: icmp_seq=35 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=36 ident=12259 ttl=62 time=60004 ms
- **64 bytes from fc00:1::3: icmp_seq=37 ident=12259 ttl=62 time=360184 ms**
- 64 bytes from fc00:1::3: icmp_seq=38 ident=12259 ttl=62 time=359160 ms
- 64 bytes from fc00:1::3: icmp_seq=39 ident=12259 ttl=62 time=358137 ms
- 64 bytes from fc00:1::3: icmp_seq=40 ident=12259 ttl=62 time=357113 ms
- ...
- 64 bytes from fc00:1::3: icmp_seq=317 ident=12259 ttl=62 time=74110 ms
- 64 bytes from fc00:1::3: icmp_seq=318 ident=12259 ttl=62 time=73087 ms
- 64 bytes from fc00:1::3: icmp_seq=319 ident=12259 ttl=62 time=72065 ms
- ...
- 64 bytes from fc00:1::3: icmp_seq=398 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=399 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=400 ident=12259 ttl=62 time=60004 ms
- --- fc00:1::3 ping statistics ---
- 400 packets transmitted, 400 received, **0% packet loss**, time 407896ms
- rtt min/avg/max/mdev = 60003.980/170727.305/360183.665/99770.604 ms, p

Fig. 6. Simulation of IP packet buffering using ping and 3 nodes

Therefore, it is relatively easy to buffer IP packets for longer periods than is typical on the Internet. It is very important to note that this buffering behavior is only needed for the forwarding nodes that have intermittent links and implement IP forwarding. Every other IP forwarder or IP node does not necessarily need to implement that behavior. For example, IP forwarders and routers on the surface of the Moon and Mars that are well-connected via 5G/WiFi/Ethernet do not need to implement this atypical IP packet buffering. Only the forwarders on deep space hops or at deep space edge are required to implement such buffering. This type of incremental deployability has been beneficial towards the success of other Internet Protocol enhancements (e.g. ECN, etc.).

Alternatively, note that if deep space forwarding nodes operate at layer 2, then the storage of layer 2 frames is done below IP, and IP is unaware of this, therefore no IP node or forwarder in the entire deep space network needs to be changed. This other approach using buffering at layer 2 can also be mixed across the network (with some nodes buffering IP packets and others buffering layer 2 frames), without any compatibility issues. This provides some design leeway to support buffering in whatever hardware or software systems can most naturally and easily accommodate in a given system. From a management perspective, the timed event or schedule-based trigger for buffering to begin or end for particular destinations can be similarly orchestrated for either layer-2 or layer-3 buffering.

B. Transport: QUIC

QUIC transport works in deep space if the stack is properly configured for the expected RTT. In particular, in a simulation, we have demonstrated that by configuring the initial_rtt and the max_idle_timeout parameters to be set to at least the expected RTT at connection establishment and by configuring the traditional congestion control to be muted and only flow control is used, then QUIC transport works by providing reliable and efficient transport, even with packet loss, reordering, or duplication.

1) HTTP Request to Voyager: In the simulation testbed, a demonstration is made sending an HTTP request to a simulated Voyager. The response is successfully demonstrated [3] even though the one-way delay is set to 18 hours, as shown in Figure 7

	Time	Source	Destination	Protocol	Length	Info
1	0.000000	192.168.65.33	192.168.65.25	QUIC	1242	Initial, DCID=d61b8e047f
2	64800.438656	192.168.65.25	192.168.65.33	QUIC	1380	Handshake, DCID=2f26ef8a
3	129600.8077…	192.168.65.33	192.168.65.25	QUIC	1242	Handshake, DCID=bf92a7a2
4	129600.8086…	192.168.65.33	192.168.65.25	QUIC	200	Protected Payload (KP0),
5	194401.1215…	192.168.65.25	192.168.65.33	QUIC	691	Protected Payload (KP0)
6	259201.4231…	192.168.65.33	192.168.65.25	QUIC	79	Protected Payload (KP0),
7	259201.4236…	192.168.65.33	192.168.65.25	QUIC	96	Protected Payload (KP0),
8	259201.4245…	192.168.65.33	192.168.65.25	QUIC	86	Protected Payload (KP0),

Fig. 7. Wireshark trace of HTTP Transaction with 18 hours delay

Various conditions such as packet reordering, duplication, and loss have also been simulated. The setup uses the Quinn [21] QUIC stack HTTP client and server where the initial_rtt and the max_idle_timeout are configured to 1.1 * RTT on connection establishment. A simpler congestion control driver is used that only manages the flow control. The delay is induced by setting the delay argument to 18 hours in the Linux netem utility. The client and server run Ubuntu 22.04. Another successful simulation takes this further by having a 5-day one-way delay (10 days RTT).

2) Quinn Workbench: The Quinn Workbench [22] is also used for the QUIC simulations in deep space. This open source software simulates a complex network of nodes and links, with delays, packet loss, duplication, and reordering. By implementing time warping, the results of a simulation over multiple days are available in seconds.

3) UDP: As noted previously, UDP is suitable for deep space use. Various protocols such as SNMP, NTP, SIP, RTP, and RTSP operate on the Internet using UDP, and are candidates for use in deep space.

SNMP, the traditional network management framework, is tested in the same simulation environment over a network with a 5-hour one-way delay and worked as expected, since the protocol itself has no notion of time. The NetSNMP [18] toolkit is used for the simulation. An SNMP get is issued with the CLI client and the timeout is set to 1.1 * the expected RTT (5 hours * 2 = 10 hours * 1.1 = 11 hours). The server is the standard NetSNMP without any configuration changes. The response is sent by the server and the client receives the response to the GET request. Therefore, to use SNMP in deep space, the requester has only to set the timeout to be larger than the expected round-trip time.

C. Applications

In the simulation testbed, we have tested various HTTP over QUIC requests such as GET, PUT and POST over large delays configured as discussed before. The client is sending those requests, and the server is configured to respond with various response sizes. The simulated network is also put under various network conditions such as packet loss, re-ordering, and duplicates using the corresponding Linux netem functionalities, and the client reliably receives the answers.

D. Network Services

1) Network Management: A key requirement for deploying a network is its management. As discussed in Section VI-B3, SNMP can be used to manage the entire network. Its recent replacement, NETCONF, is typically used over TCP on the Internet but there is an HTTP-based variant, RESTCONF, that is viable for deep space use over QUIC. Moreover, NETCONF natively over QUIC is also in the process of being standardized [7] and is another option.

2) DNS: In the initial phases of deployment, deep space IP network nodes can be identified by their IP addresses. However, soon the scalability and operational impacts of using IP addresses will require using naming. Domain Name System (DNS) can be used in deep space if properly deployed [6]. It enables large-scale deployment and distribution of names in a secure and authenticated way.

VII. Deep Space IP Stack

Fig. 8 shows the resulting deep-space IP stack described in this paper.

Although the use of the IP stack has been demonstrated for deep-space networking, it should be emphasized that the typical Internet applications and usage known and used daily on Internet are typically not appropriate for deep space. The IP stack enables space-specific applications to be used, not the normal ones on the Internet. Given that Moon is just a few seconds away, we may see slightly tuned Internet applications

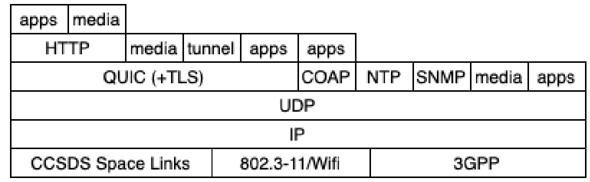

Fig. 8. Deep space IP stack

to be usable for Moon when there is no interruptions. However, since communication interruptions are also expected for the Moon, all the techniques described in this paper are also required for Moon deployment.

A. Advantages

Compared to the Bundle Protocol, the use of IP stacks in deep space opens up a large toolset of mature protocols, applications, frameworks, security, software and knowledge to be reused for space. This significantly reduces the costs and risks associated with supporting the growing array of complex deep space missions. These missions require networking that is reliable, secure, and dependably managed, and find value in leveraging IP technology to meet their needs, instead of relying on DTN BP. This approach reuses software that has been well-tested and exercised at large scale, and under other stressing conditions in the Internet. Security protocols for the Internet have been scrutinized heavily and hardened specifically for needs of compliance with commercial, medical, government, and defense/intelligence networking requirements. High-bandwidth IP processing is implemented in different forms of hardware acceleration in terrestrial systems and can also be applied for deep space networking with proper environmental hardening.

B. Standards

The Internet Engineering Task Force created the ”Taking IP to Other Planets”(tiptop) [29] working group to define the profiles and adaptations for the Internet Protocol stack to be used in deep space, as described in this paper. Standards set in this working group enables interoperability between implementations and lower the costs of implementation, procurement and operations for space users.

VIII. Conclusion

This paper demonstrates the use of the Internet Protocol stack for deep space, implementing, in fact, a delay- and disruption-tolerant network, contradicting prior work that ruled out this possibility. It is based on the following adaptations: buffering packets for IP forwarding facing interruptions, properly configuring the QUIC transport stack using the expected RTT, configuring timeouts of applications and designing them with the asynchronous pattern. The capability of IP in deep space has been demonstrated through various simulations.

ACKNOWLEDGMENT

We acknowledge the following individuals for their contribution during this work: Jean-Philippe Dionne, Adolfo Ochagavia, Christian Huitema, Martin Thompson, Benjamin Saunders, François Michel, Maxime Piraud, Martin Duke, Lars Eggert, Dean Bogdanovic, Tony Li, Carlos GomezMontenegro, Émile Stephan, Warren Kumari, Mark Andrews, Vint Cerf, Yosuke Kaneko, Atsushi Tagami, James Schier, Felix Flentge, and Juan Fraire. There have also been countless emails and hall discussions with additional people that we cannot list here, but are acknowledged.

REFERENCES

[1] D. Belson and L. Purdue, “Examining HTTP/3 usage one year on“,June 6th 2023, https://blog.cloudflare.com/http3-usage-one-year-on.
[2] “IP Protocol Stack for Deep Space“, https://deepspaceip.github.io/
[3] “Deep Space IP Simulation Testbed Results“, https://deepspaceip.github.io/testbed
[4] M. Blanchet, C. Huitema and D. Bogdanović, “Revisiting the Use of the IP Protocol Stack in Deep Space: Assessment and Possible Solutions“, March 2024, https://datatracker.ietf.org/doc/draft-many-deepspace-ip-assessment/
[5] M. Blanchet, “Bundle Protocol Issues” Internet-Draft, draft-any-dtn-bp-issues, IETF, September 2025.
[6] Marc Blanchet, “Deployment and Use of the Domain Name System (DNS) in Deep Space,” Internet-Draft draft-many-tiptop-dns-00, IETF, 1 July 2025. Available: https://datatracker.ietf.org/doc/draft-many-tiptop-dns/
[7] Jinyou Dai, Shaohua Yu, Weiqiang Cheng, Marc Blanchet, and Per Andersson, “NETCONF over QUIC,” Internet-Draft draft-ietf-netconf-over-quic-04, IETF, 22 May 2025. Available: https://datatracker.ietf.org/doc/draft-ietf-netconf-over-quic/
[8] Marc Blanchet, Wesley Eddy, and Tony Li, “An Architecture for IP in Deep Space,” Internet-Draft draft-many-tiptop-ip-architecture-01, IETF, 7 July 2025. Available: https://datatracker.ietf.org/doc/draft-many-tiptop-ip-architecture/
[9] Marc Blanchet, Wesley Eddy, and Marshall Eubanks, “IP in Deep Space: Key Characteristics, Use Cases and Requirements,” Internet-Draft, draft-ietf-tiptop-usecase-00, IETF, 20 July 2025. Available: https://datatracker.ietf.org/doc/draft-ietf-tiptop-usecase/
[10] Marc Blanchet, “QUIC Profile for Deep Space,” Internet-Draft draft-many-tiptop-quic-profile-01, IETF, 23 August 2025. Available: https://datatracker.ietf.org/doc/draft-many-tiptop-quic-profile/
[11] “The Future Mars Communications Architecture“, Volume 1, Report of the Interagency Operations Advisory Group Mars and Beyond Communications Architecture Working Group, 2022, https://www.ioag.org/Public
[12] “The Future Lunar Communications Architecture“, Volume 1.3, Report of the Interagency Operations Advisory Group Lunar Communications Architecture Working Group, 2022, https://www.ioag.org/Public
[13] Grant Ryden and Michael Volle, “NASA Lunar Communications Relay and Navigation Systems (LCRNS) Reference Constellation 3.1,” NASA White Paper, April 3, 2025.
National Aeronautics and Space Administration. Document ID 20250002698. Available at: https://ntrs.nasa.gov/citations/20250002698
[14] Maxim Krasnyansky, “Universal TUN/TAP device driver“, https://docs.kernel.org/networking/tuntap.html
[15] Gladden, Roy E., “Mars Relay Operations Service (MaROS): a present service preparing for the future“, https://hdl.handle.net/2014/45530, SpaceOps 2014 13th International Conference on Space Operations, Pasadena, California, May 5-9, 2014
[16] Marc Blanchet, ”Earth-Mars Communication Windows Usage Study”, Deepspace IP informal meeting, IETF 121, Dublin, 4 November 2024. Available: https://deepspaceip.github.io/meetings/ietf121/ietf121-deepspace-mars-comm-windows.pdf
[17] “NASA’s Mars Fleet Will Still Conduct Science While Lying Low“, https://www.jpl.nasa.gov/news/nasas-mars-fleet-will-still-conduct-science-while-lying-low
[18] “NetSNMP“, http://www.net-snmp.org/
[19] Nokia Bell Labs, “Space Communications,” Network Fundamentals, Bell Labs, Nokia. Available: https://www.nokia.com/bell-labs/research/network-fundamentals/space-communications/
[20] “Deep Space Optical Communications (DSOC)“, NASA Jet Propulsion Laboratory, https://www.jpl.nasa.gov/press-kits/psyche/dsoc
[21] “Quinn‘, https://github.com/quinn-rs/quinn
[22] “Quinn Workbench‘, https://github.com/aochagavia/quinn-workbench
[23] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and H. Weiss, “Delay-Tolerant Networking Architecture“, RFC 4838, DOI 10.17487/RFC4838, April 2007, ¡https://www.rfc-editor.org/info/rfc4838¿.
[24] Iyengar, J., Ed., and M. Thomson, Ed., ”QUIC: A UDP-Based Multiplexed and Secure Transport”, RFC 9000, DOI 10.17487/RFC9000, May 2021, ¡https://www.rfc-editor.org/info/rfc9000¿
[25] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., “HTTP Semantics“, STD 97, RFC 9110, DOI 10.17487/RFC9110, June 2022, ¡https://www.rfc-editor.org/info/rfc9110¿.
[26] Mike Bishop, “HTTP/3,” RFC 9114, Internet Engineering Task Force, June 2022. Available at: https://datatracker.ietf.org/doc/rfc9114/
[27] Burleigh, S., Fall, K., Birrane, E., and III, “Bundle Protocol Version 7“, RFC 9171, DOI 10.17487/RFC9171, January 2022, ¡https://www.rfc-editor.org/info/rfc9171¿.
[28] Marc Blanchet, “Using the Internet Protocol Suite for Deep Space Networking”, Proceedings of the SpaceOps 2025 Conference, May 2025.
[29] “Taking IP To Other Planets (tiptop)‘, https://datatracker.ietf.org/group/tiptop/about/