

# Using the Internet Protocol Suite for Deep Space Networking

Marc Blanchet  
Viagenie  
Québec, Canada  
marc.blanchet@viagenie.ca

**Abstract**—This paper describes how to use the Internet Protocol suite to implement a network in deep space. It lists the key adaptations that are needed to make it work, such as temporary storage of IP packets on forwarders with intermittent links, configuration of the QUIC transport protocol with parameters related to the expected round-trip time and by setting application timeouts also related to the expected round-trip time. Simulations described in this paper confirmed the suitability of the IP suite in deep space.

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

## I. INTRODUCTION

Deep space has been solely conquered by government agencies such as NASA, ROSCOSMOS, ESA, JAXA, CNES, CSA and others, mainly due to the very high costs and risks associated with space missions. Communications are largely direct between terrestrial infrastructures and spacecrafts. In networking jargon, these are point-to-point communications. As shown in Fig. 1, the only exception is the use of relays often in orbit around a planet like Mars, essentially orbiters or satellites, allowing the retransmission of communications to and from less powerful devices on the planet's surface, such as exploration robots on Mars.




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. Long interruptions can be from a few seconds to several days or weeks. For example, Mars may be on the other side of the Sun from Earth, as illustrated in Fig. 2, creating a period where direct communications are not possible, a situation that occurs every 2 years. A recent Martian solar conjunction took place in November 2023 when, for two weeks, no direct communication was possible between the two planets[MARSCONJUNCTION].



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

## II. RELAYS

The relays are currently relatively unintelligent in the sense that they do not create a true network because, for example, there is no consideration of alternate paths nor do they make any routing decisions. However, during time slots where no communication is possible, for example when the Mars relay is on the opposite side of the planet to the rover, the relay temporarily stores the frames until communication is re-established. For Mars, communication delays and interruptions are planned and managed by a brokering software [MAROS] where missions such as the Deep Space Network and relays around Mars provide communication windows to other missions such as those of exploration robots or rovers on the surface of Mars. Everything is specified and planned in advance. 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.

### A. Mars Communication Windows and Interruptions

Currently, on Mars, there are multiple orbiters acting as relays: Mars Reconnaissance Orbiter(MRO), Odyssey(ODY), Mars Atmosphere and Volatile Evolution(MVN, Maven), Trace Gas Orbiter(TGO) and Mars Express(MEX). There are four spacecrafts on the Mars surface: Mars Science Laboratory(Curiosity), Mars Perseverance Rover (M20) and In-sight(NSY). The latter is currently unreachable and the mission may soon be declared terminated. The helicopter Ingenuity is also there, but is not directly used in communications to Earth, as its communications is proxied by the Curiosity rover. The MAROS project enables the matching of communication windows for the purpose of the various missions.

## B. Mars Round-Trip and Holding Times

The list of potential communication windows and actual occurred ones between the Mars orbiters and the rovers, and between the orbiters and the antennas on Earth, were provided by JPL. The data spans over a period of April to June 2024. A complete path is from Earth to orbiter to rover to the same orbiter to Earth. All the possible complete paths were calculated by matching the time and sequence of periods, with disjoint periods. As shown in table I, out of 4564 communication windows between Mars orbiters and rovers, 809 complete paths were computed as usable paths. The average complete path time is 12 hours and 40 minutes, the maximum is 161 hours and 42 minutes 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

These are theoretical as most communication windows between orbiters and rovers are typically not used by the missions, for various reasons, such as the rover is moving, the rover does not have enough power, etc. These are also not optimal as they were calculated as disjoint, where in fact, if windows overlap, then communications can be faster in total time. These computed numbers are very important for this work since it shows the large range of round-trip times and the holding times of data in orbiters that can occur in deep space because of the intermittence of communications through orbiters. This is an important input into further characterization of the communications in space and how to properly engineer networking given that data.

## C. Moon Case

Since there are currently no communicating assets on Moon, similar data can not be used at this time. However, while Moon is just a few seconds away from Earth, the intermittence of communications through orbiters will also happen. In both cases, if the coverage around the celestial body is complete with a constellation of satellites/orbiters, then the intermittence may disappear. However, this will take quite a long time to accomplish given the costs of sending spacecrafts to space.

Those Mars results have been used in the simulations. If the IP stack and applications works with those numbers, then the Moon case will also work as it is a simpler problem.

## III. DECREASED COSTS, FASTER COMMUNICATIONS AND COMMERCIALIZATION

Nowadays, space conquest has undergone significant transformations. Indeed, industrialization now allows for a decrease in the construction costs of space devices. A good example

is the Cubesat [5], which consists of a simple and relatively small format of satellites, allowing less affluent actors and projects to send devices into space at a lower cost. Furthermore, the space industry now uses less demanding electronic components [4], allowing for significant cost reductions but requiring more verification tests, including radiation testing. Finally, launcher innovation led by SpaceX has significantly reduced launch costs. Finally, recent tests [14] [19] show that optical communications using lasers will be possible in deep space, enabling much larger and symmetrical bandwidth. These combined changes create a new era where many small to large companies and organizations become players in this space conquest. Space agencies used to define the specifications themselves and then subcontract the realization are now in a mode of contracting for commercial services where they can rent services according to their needs. These services include communications and networking. In some ways, the Internet has seen similar history where the initial funding was from governments and then commercialization and openness enables a large network to be built creating a platform for innovation. The key to make this happen is to use standards, ensuring interoperability between all protocols, software and infrastructure.

## IV. MOON AND MARS NETWORKING ARCHITECTURE

[11] and [10] describe the communications and network architecture envisioned for Moon and Mars. On the surface of the celestial bodies and around such as cislunar, the networking stack will use the Internet Protocol(IP) suite over the Earth wireless technologies such as 5G/6G and Wifi. Therefore, by using IP over CCSDS deep space links, a single networking stack over the whole path is used, enabling faster delivery, reliability, routing, and cost-effective platforms.

## V. BUNDLE PROTOCOL

In early 2000s, Vint Cerf and JPL engineers investigated if the Internet Protocol(IP) suite would be usable to create a network in deep space [RFC4838], named the Solar System Internet [12]. We use the term deep space here to clearly differentiate with the various orbits (LEO, MEO, GEO) around Earth where IP is already in use. Deep space here means communications with the Moon and its orbits, Mars and beyond as well as cruising spacecrafts.

Their conclusion was that the IP protocol stack is not usable in deep space given its unique characteristics, such as very large and variable delays and long communication interruptions, where the current IP transport and applications at that time were too chatty for that use case. A network protocol working with long delays and interruptions is named Delay and Disruption Tolerant Networking (DTN). It was deemed necessary to create a completely new network stack using a store and forward paradigm. The result is the Bundle Protocol [25], which is different and incompatible with IP. Since then, there has been work in both the Internet Engineering Task Force [9] and the Consultative Committee for Space Data Systems [3] to engineer all the pieces necessary for a full

functional networking stack, obviously a major undertaking, which is still not quite complete. Over time, various issues were found with BP [1], such as: no end to end reliability, no route aggregation, non scalable routing, etc.

## VI. RESASSESSMENT OF IP FOR DEEP SPACE

We reassessed [8] the use of the IP stack in deep space. Since the initial assessment in early 2000's, the IP stack has evolved a lot, including the capability of handling Internet of Things (IoT) which have many similarities with space usage, such as low bandwidth, low energy, intermittent communications, low cpu and low memory. Since early 2000's, the QUIC transport protocol [23] has been invented and is now a significant portion of Internet traffic [2].

The primary conclusion of this new assessment is that the IP stack is usable in deep space, for the following reasons.

### A. Network Layer

At the networking layer, an Internet Protocol packet does not carry any notion of time. IPv4 does have a Time-To-Live(TTL) field, however, while it was initially designed as decreasing time by intermediate nodes, it has been implemented as a hop count, which was then officialized in IPv6 as the Hop Count field. Therefore, an IP packet by itself can live for a very long time without any issue. At the IP layer, the only expiration is when the maximum number of hops is reached.

### B. Transport

At the transport layer, as noted in [RFC4838], TCP is too chatty with the 3-way handshake and with the defined timers. One should also add the TLS handshake to the budget, so TCP is not suitable for deep space. However, UDP has no notion of time, therefore can be used as is in deep space. While QUIC was not designed for space use, it can be profiled to support that use case. QUIC is a modern transport and has many features very useful for space. It uses a single round-trip time(RTT) for establishing a secure and encrypted connection between two peers, while TCP and TLS require multiple RTTs. The connection can carry an unlimited number of requests and answers, named streams, without having to re-establish the connection itself, while TCP would require multiple connections. Any peer can start a new stream. It supports the re-establishment of a connection with data in the payload of the first packet, a feature named ORTT. There is no a priori expiration of a connection, so it can remain opened for minutes, hours, days, weeks, months or even years which is well suited for space. A deployment scenario could be to even establish the connection before the spacecraft leaves Earth and keep the connection for the whole duration of the mission.

If the spacecraft changes its network attachment, as when a cruising spacecraft is landing on a planet and attach to the surface network, the spacecraft will receive a new IP address or range. QUIC solves the mobility problem because there is a security association between the two peers and because of its independence on the IP addresses and ports to identify the connection, since it has its own connection identifier within

the protocol. The connection is re-established even when the IP address or port changes, supporting full mobility.

As a reliable transport, QUIC automatically manages issues in networks such as packet reordering, duplication and loss and pacing.

However, the typical QUIC behaviour for Internet assumes short delays in the order of hundred milliseconds and relative full time connectivity. Therefore, it does not work as is for deep space. However, by setting a few of its transport configuration parameters and implementing a simpler congestion control mechanism, it works well in deep space conditions.

### 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, current Internet applications will most likely not work as is in deep space. However, if the applications are designed with the expected large RTT and with using the asynchronous pattern, then they become suitable for deep space.

The most used application protocol nowadays is HTTP [24]. 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. Therefore, HTTP can be used in deep space. The typical transport of HTTP is TCP which, as noted above is not suitable for deep space. However, HTTP is also defined and used over QUIC for a large part of the HTTP traffic on the Internet. Therefore HTTP over QUIC is required.

Application and application protocols using UDP transport, such as streaming, are also suitable for deep space, given proper timeout configuration.

## VII. ARCHITECTURE AND SIMULATION RESULTS

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

### A. Forwarding

As discussed previously, deep space communications are interrupted regularly because the relays cannot always provide a complete path from source to destination.

On Internet, an IP forwarder (or router) drops packets when the destination is seen unreachable in its forwarding table and therefore sends an ICMP destination unreachable error message to the source. However, in space, given that orbital relays will have regular disconnections, the packets should not be dropped but instead stored temporarily until the communication is re-established.

This storage behaviour can be accomplished at a layer below IP, as done currently by the Mars orbiters. If the relays are doing the storage at layer 2, then the IP layer is unaware of that storage and no modification to IP behaviour is necessary.

If an IP forwarding node has links with intermittent windows of communications, then that node behaviour should

be modified. Instead of dropping packets, as noted earlier, it should instead store the packets until the next window of communication to the next hop appears, when it will then resume forwarding to the next hop.

Policies can be established for storage or forwarding based on various criteria, such as source or destination addresses or prefixes, diffserv/traffic class field, flow label or any other. Tools to implement those policies are already available in all platforms.

The temporary storage of IP packets was implemented with 200 lines of C code, in a prototype using the Linux Tun interface [15] and a simulation showed [7] that it was able to store packets during a link interruption and that packets were then pulled from storage and forwarded when the link becomes up again, using a simple FIFO policy. Ping was used for testing. The network consisted of three nodes: a client node, a server node and a forwarder node, as shown in Fig. 3.

A delay of 30 seconds was set on the forwarder node on both directions. At the beginning of the simulation, only the induced delay was involved, therefore the ping client receives the replies after 60 seconds. The forwarder node then have an interruption of its link to the server node for a period of 5 minutes. The ping client is not receiving any reply or any error during that period of time. In a normal Internet forwarder environment, the forwarder would have sent ICMP destination unreachable error messages to the ping client. However, in this implementation for deep space simulation, the forwarder stores the packets. After the 5 minutes period, the link is back and the forwarding resumes. The ping client starts receiving replies with a 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. It should be noted 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.

Therefore, it is relatively easy to store IP packets temporarily for longer periods than typically on Internet. It is very important to note that this storage behaviour is only needed for the forwarding nodes that have intermittent links and are implementing IP forwarding. Every other IP forwarder or IP node should not implement that storage behaviour. For example, IP forwarders and routers on the surface of the Moon and Mars running 5G/Wifi/Ethernet would not need to implement storage. Only the forwarders in space or at space edge are required to implement storage. If the forwarding nodes are L2, then the storage of frames is done at that layer and IP is unaware of storage, therefore no IP node or forwarder in the whole deep space network has to be changed.

### B. Transport

As noted in [22], the TCP transport is not suitable for deep space. QUIC transport works 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

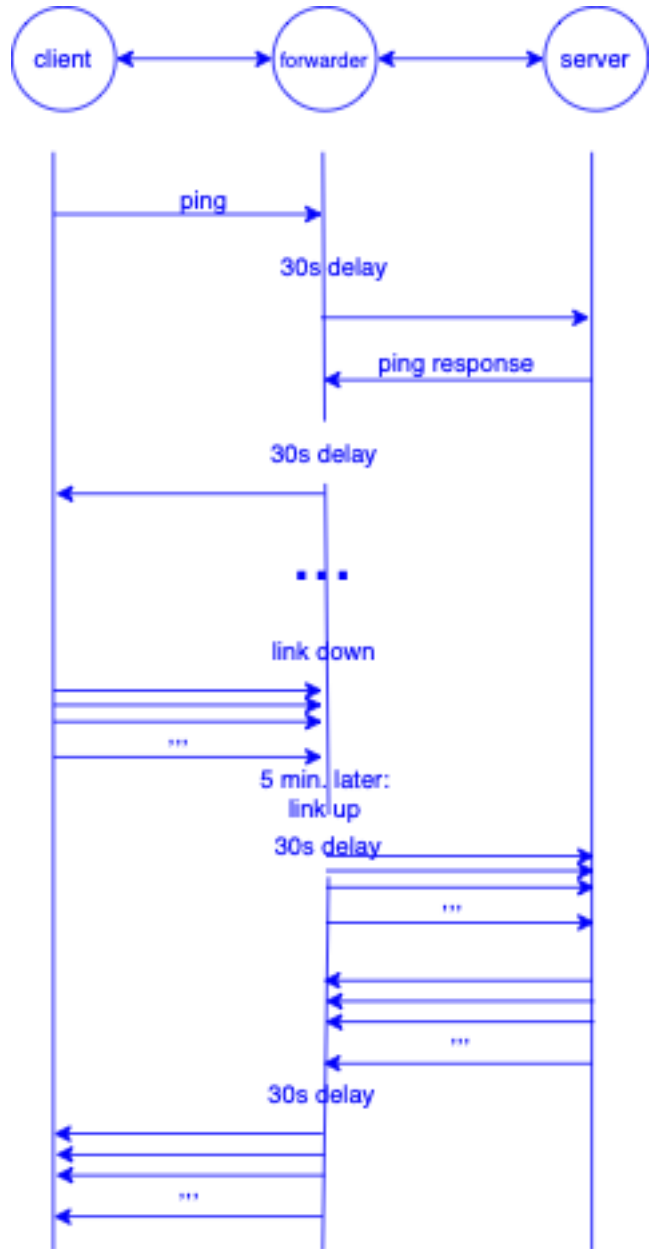


Fig. 3. Simulation of IP packet storage using ping and 3 nodes

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 does work by providing reliable and efficient transport, even with packet loss, reordering or duplication.

In the simulation testbed, sending an HTTP request to Voyager and getting the response was successfully demonstrated [7]. Not that the real Voyager received the packet, but the one-way delay in the simulation was set to 18 hours. Various conditions such as packet reordering, duplication and loss have also been simulated. The setup uses the Quinn [20] QUIC stack HTTP client and server where the initial\_rtt and the

max\_idle\_timeout were configured to  $1.1 * RTT$  on connection establishment. A simpler congestion control driver was used that only manages the flow control. The delay was induced by setting the delay argument to 18 hours of the Linux netem utility. Client and server were running Ubuntu 22.04. Another successful simulation took this further by having a 5 day one-way delay (10 days RTT).

The Quinn Workbench [21] was used for QUIC simulations. This open-source software simulate 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.

As noted previously, another transport protocol, User Datagram Protocol(UDP), works just fine in deep space, since it has no notion of time. Various protocols such as SNMP, NTP, SIP, RTP, RTSP are using UDP, therefore are candidate to work in deep space. SNMP, the traditional network management framework, was tested in the same simulation environment over a network with 5 hours one way delay and worked just fine, since the protocol itself has no notion of time. The NetSNMP [18] toolkit was used for the simulation. An SNMP get was issued with the CLI client and the timeout was set to the  $1.1 * \text{the expected RTT}$  (5 hours \* 2 = 10 hours \* 1.1 = 11 hours). The server was the standard NetSNMP without any configuration change. The reply was sent by the server and the client received the answer to the GET request. Therefore, to use SNMP in deep space, the requester has only have to set the timeout to be larger than the expected round-trip time.

### C. Applications

The third important aspect of IP in deep space rely on the application layer itself. If the application uses HTTP, the HTTP protocol itself does not have any notion of time. Only clients, servers and applications usually have timeouts that should be adjusted to values of the expected round-trip time. Any application in deep space must be designed completely asynchronously (as it is also the best current practice for Internet applications nowadays). It must also contain either no timeouts or have timeout set properly related to the expected round-trip time. In the simulation testbed, we have tested various HTTP requests such as GET, PUT and POST over large delays over QUIC configured as discussed before. The client was sending those requests and the server was configured to answer with various response sizes. The simulated network was also put into various network conditions such as packet loss, reordering and duplicates using the corresponding Linux netem functionalities and the client received reliably the answers.

## VIII. DEEP SPACE IP STACK

Fig. 4 shows the resulting deep space IP stack.

While 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 the space specific applications to be used, not the




Fig. 4. Deep space IP stack

normal Internet ones. However, given that Moon is just a few seconds away, we may see slightly tuned Internet applications to be usable for Moon when there is no interruptions. However, as communication interruptions are also expected for Moon, then all the techniques described in this paper are needed also for Moon deployment.

### A. Advantages

Compared to the Bundle Protocol, the use of IP stack in deep space opens up a large toolset of protocols, applications, frameworks, security, software and knowledge to be reused for space. It significantly decreases the communication and networking costs of space missions, compared to the special Bundle Protocol stack. It also decreases the risk as instead of using specialized space software as today, it reuses software that has been exercised at large scale. Security protocols for Internet have been scrutinized as no other therefore decreasing the security risk for space communications. High-bandwidth hardware acceleration is available for Internet use and can also be used for deep space networking with proper space hardening.

The Bundle Protocol suite does not provide any transport semantics to recover from issues such as bundle reordering, bundle duplication and pacing. It does provide an optional bundle loss recovery mechanism but only on a hop by hop function. Using a reliable transport such as QUIC over IP provides all those semantics and guarantees delivery of data end to end.

### B. Standards

The Internet Engineering Task Force started the Taking IP to Other Planets(tiptop) [26] 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 will enable interoperability between implementations and lower the costs of implementation and procurement for space users.

## IX. CONCLUSION

This paper demonstrates the use of the Internet Protocol stack for deep space, implementing in fact a delay and disruption-tolerant network. It is based on the following adaptations: storing 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 would like to 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, Wesley Eddy, Carlos GomezMontenegro, Émile Stephan, Warren Kumary, Mark Andrews, Vint Cerf, Yosuke Kaneko, Atsushi Tagami, James Schier, Felix Flentge, Juan Fraire. There have also been countless emails and hall discussions with additional people that we cannot list here, but are acknowledged.

## REFERENCES

- [1] M. Blanchet, “Deep Space Networking: Bundle Protocol Assessment“, April 2025, <https://datatracker.ietf.org/doc/draft-many-dtn-bp-assessment/>
- [2] D. Belson and L. Purdue, “Examining HTTP/3 usage one year on“, June 6th 2023, <https://blog.cloudflare.com/http3-usage-one-year-on>.
- [3] “Consultative Committee for Space Data Systems (CCSDS)“, <https://ccsds.org>
- [4] P. Majewicz, “NASA Efforts In Utilizing Commercial-Off-The-Shelf (COTS) Electronics In Mission Systems“, , 2022, <https://ntrs.nasa.gov/api/citations/20220015267/downloads/20220015267-Majewicz-ACCEDE-2022-Presentation-v3.pdf>
- [5] CubeSat, “Origin of the New Space Revolution“, <https://www.cubesat.org>
- [6] “IP Protocol Stack for Deep Space“, <https://deepspaceip.github.io/>
- [7] “Deep Space IP Simulation Testbed Results“, <https://deepspaceip.github.io/testbed>
- [8] 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/>
- [9] “Delay/Disruption Tolerant Networking (dtn)“, Internet Engineering Task Force, <https://datatracker.ietf.org/wg/dtn/about/>
- [10] “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>
- [11] “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>
- [12] Internet Society Interplanetary Chapter, <https://ipnsig.org>
- [13] Internet Society, <https://isoc.org>
- [14] “Laser Communications Relay Demonstration (LCRD) Overview“, NASA <https://www.nasa.gov/directorates/stmd/tech-demo-missions-program/laser-communications-relay-demonstration-lcrd-overview/>
- [15] Maxim Krasnyansky, “Universal TUN/TAP device driver“, <https://docs.kernel.org/networking/tuntap.html>
- [16] 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
- [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] “Deep Space Optical Communications (DSOC)“, NASA Jet Propulsion Laboratory, <https://www.jpl.nasa.gov/press-kits/psyche/dsoc>
- [20] “Quinn“, <https://github.com/quinn-rs/quinn>
- [21] “Quinn Workbench“, <https://github.com/aochagavia/quinn-workbench>
- [22] 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>.
- [23] 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>
- [24] 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>.
- [25] 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>.
- [26] “Taking IP To Other Planets (tiptop)“, <https://datatracker.ietf.org/group/typtop/about/>