

Improving the Performance of Asynchronous Communication in Long-Distance Delay-Sensitive Networks through the Development of Context-Aware Architectures

Cathryn Peoples, Gerard Parr, Bryan Scotney,
Adrian Moore
University of Ulster, School of Computing and
Information Engineering, Coleraine,
Northern Ireland, BT52 1SA.
{peoples-c1; gp.parr; bw.scotney;
aa.moore}@ulster.ac.uk

Petre Dini
Cisco Systems
San Jose
California
pdini@cisco.com

Abstract

Context-awareness is inherent in anticipated interplanetary missions. Swarm technologies [1] use context-awareness in short-haul networks between components, and long-distance networks allow communication with Earth. However, the propagation delays limit real-time communications, deep space being an environment in which the speed of light becomes a restriction. Therefore, the development of a protocol stack which is adaptive to application requirements and external influences will help to maximise communication synchronicity.

As part of a first year doctorate research programme, this paper correlates current stack functionalities with interplanetary application requirements. A redesigned stack proposes to resolve this misalignment. Context-awareness is incorporated, enabling intelligent protocol selection using application layer knowledge and environmental information, with particular attention given to transport protocols. The paper concludes by considering transport protocol characteristics when deployed beside a context-aware layer, with the long-term aim being the development of a transport protocol suitable for deployment in the state-of-the-art context-aware stack.

Keywords: *Context-awareness, delay-sensitive protocols, Quality of Service (QoS), Interplanetary networks.*

1. Introduction

While working groups have been formed for this purpose, network protocol stacks proposed for the interplanetary environment remain incomplete. Although these stacks can cope with Earth Observing (EO) missions and exploits to Mars, they lack features required for envisaged missions. To achieve the vision of space exploration [1], radical networking approaches are needed. The intention is to deploy a permanent infrastructure, at least as far as Mars, in the near future. However, this will only become a reality when communication techniques which maximise hardware utility are developed to justify the costs involved.

Interplanetary network infrastructures were originally built and removed for each mission [2]. Components below the Tracking Delay Relay Satellite System (TDRSS) are permanent, while spacecraft are temporary. Network endpoints therefore change and permanent paths cannot be assumed. Also, all communication is managed by mission control on Earth, limiting any ability of responding to unexpected events, to capture unpredicted scientific occurrences or cope with mission failure. Given the state-of-the-art processing capabilities on Earth, the restricted developments in space are obvious.

Initial experiments into automating communication in deep space include the Communications and Navigation Demonstration on Shuttle (CANDOS) mission [3]. While Columbia, the spacecraft carrying the experiment, disintegrated on re-entry, the telecommunication aspect of the mission was

considered to be successful [3]. Mobile IP automated end-to-end connectivity, and the User Datagram Protocol (UDP) [4] and Multicast Dissemination Protocol (MDP) [5] enabled bidirectional communication. This mission was the first time that a file travelled to Earth from space without having its communication path previously defined [6].

CANDOS displayed the ability to extend the Internet protocols into space. More recently, Cisco sent a Mobile Access Router into space as part of the Cisco Low Earth Orbit (CLEO) project [7], proving that commercial IP technology is suitable for missions. Located on-board spacecraft or satellites, routers enable a more intelligent transporter of data signals between Earth and space. However, the harsh environment in outer space causes hardware and software implementation issues. With regard to software, deficiencies in the performance of network protocols over long-distances become obvious. Current missions to Mars are controlled by a stack from the Consultative Committee of Space and Data Systems (CCSDS) [8]. However, while a strong foundation has been set, further improvements are needed before manned missions to Mars become a reality.

The overall aim of interplanetary networks is to achieve autonomous communication with space as on Earth. From Cisco's point of view, a research gap exists and there is potential for them to be key players in the construction of the next generation network (NGN) in space. Securing the hardware aspect should be routine due to the amount of experience in this field. Resolving the software aspect, however, will be more challenging.

This paper attempts to further research regarding the gap between current and required capabilities of long-distance protocols: it considers the achievements of a reconstructed protocol stack with context-aware capabilities. The operational characteristics of interplanetary applications are considered, with the aim of correlating performance requirements with stack complexity. The construction of an optimal transport protocol in this context-aware stack is also considered. Such an approach has not been previously provided. The race has begun to prove its value.

2. Context-Awareness

Context-awareness is the element which we (the authors) consider to be missing from current protocol stacks. Context-awareness describes the ability to consider external factors when making operational decisions. Contexts include resource availability, location constraints, application requirements, or environmental conditions, as shown in Figure 1.




Figure 1 Types of Context-Awareness

Figure 1 builds on a model introduced by Khedr and Karmouch (2002) [9], with the incorporation of environmental awareness. Environmental awareness is important in long-distance scenarios, being the primary factor for communication difficulties. The negative effects of long propagation times are heightened by limited memory resources and intermittent connectivity. Without resource scheduling, careful protocol design, and efficient network management, transmissions can easily be lost. Context-aware capabilities can reduce dependence on the network management function and improve mission functionality.

Contexts may be sensed in the hop-to-hop or end-to-end environment. End-to-end sensing retrieves context information once for the entire link, while contexts are sensed on intermediary links in the hop-to-hop approach. The hop-by-hop approach may be more appropriate in deep space: interplanetary communication occurs over vast distances and it is expected that circumstances will change with distance. Therefore, it is more suitable to deploy a hop-by-hop approach and make decisions on recent information. Both approaches, however, increase the duration of signal propagation. While this cost may be minimised by end-to-end sensing, there is no guarantee that decisions based on context information for the whole link will remain relevant for the duration of the communication. Context-awareness must therefore be implemented with caution to ensure that the costs involved do not exceed the benefits received.

Real-time communication is difficult in the interplanetary environment, a significant issue due to the QoS requirements of typical applications. Context-awareness can relieve this problem: where real-time communication is not possible, a pseudo real-time service can be achieved through effective protocol choice when transmissions do not have mission-critical importance. For other applications, real-time communication is necessary and an inability

to overcome the propagation delays will impede the path of future exploration. However, envisaged technologies require the integration of autonomy and context awareness to allow performance of their routine tasks. With swarm technology, command and control data will be retrieved from a controlling element within the swarm [11]. Mission-critical data will therefore be communicated over short distances and communication with Earth will occur only for scientific data relay. In enabling these autonomic capabilities which respond to context information, the transmission of mission-critical information over interplanetary distances becomes unnecessary, and applications which succumb to the long-propagation delays will have non mission-critical requirements. The context-aware element is therefore important in enabling the mission-critical information to be made available and used effectively.

3. Interplanetary Applications

In determining the potential for applying context-awareness to interplanetary applications, an understanding of current performance levels must be gained. Typical interplanetary applications are characterised according to their QoS metrics in Figure 2.

Application	Physical Demands		Network Services			Coping Abilities	
	Typical Data Volume (kbps)	Store and Forward Requirements	Transmission Thresholds	Acknowledgements	Retransmission Ability	Tolerance of Poor Timing Estimations	Error Coping Abilities
Telemetry	28.8	Low	< 100 ms	On	On	Low	Low
Scientific Data	-	Medium	2 - 4 s	Off	On	High	High
Command & Control	28.8	Low	< 100 ms	On	On	Low	Low
Network Performance Monitoring		Medium	150 - 400 ms (preferably < 100 ms)	Off	On	Medium	Medium
Multimedia (Voice, Audio and Video)	4 - 700	High	100 ms	Off	Off	Low	Medium

Figure 2 Application Performance Metrics

The applications have been characterised in relation to their demands on the network, services required, and coping abilities when network conditions are sub-optimal. Physical network demands include bandwidth and memory requirements. Services include acknowledgements, retransmissions, and latency thresholds. Coping abilities with poor timing estimations and errors are also measured.

Context-awareness can maximise the possibility that the required QoS levels will be reached in interplanetary networks. Acknowledgements are

ideally switched off for all applications to maximise performance in the limited bandwidth environment and will only be switched on when the application is mission-critical or where bandwidth redundancy exists. It will be the function of the context-aware application to determine the network's ability to support acknowledgements and the advantages from their use. In the case where the application is not tolerant of poor timing estimations and bandwidth is unsuitable, the transmission may be halted where the application is not mission-critical. This will free bandwidth for mission-critical applications and those with stringent resource and latency requirements: applications without mission critical importance can transmit at a more appropriate time. Operating in this manner will improve the overall performance of the network, avoiding the consumption of resources by transmitting when it is unlikely to reach its destination with the required QoS levels. However, the overall complexity of the decision to transmit becomes obvious.

Network services can change as communication progresses. As a function of distance, it may become necessary that acknowledgements are eventually switched off as distance increases. Reduced bandwidth and the risk of errors may make it necessary to concentrate all resources on transmitting the actual data, without retransmissions and extra overhead. Network services can be applied in response to application layer knowledge and context knowledge gained as the transmission progresses.

The requirement for context-awareness is evident in envisaged deep space missions. Missions using Autonomous Nano-Technology Swarm (ANTS) technology [10] [11] become possible with self-configuring capabilities. Although not explicitly outlined in ANTS, inclusion of context-awareness is inherent – the spacecraft reconfigures and makes decisions based on information sensed.

To conclude, while intelligence will make possible missions previously impossible, context-awareness will increase processing time at the source. However, this increase in processing time may result in an overall higher performing network. Future work must therefore evaluate these parameters against each other.

4. Current Protocol Stacks

Context-awareness must be incorporated into network protocol stacks. Current stacks have not made allowances for this functionality. There are three main protocol stacks responsible for interplanetary communications, including the traditional Open Systems Interconnection model (OSI) developed by the International Standards Organisation

(ISO), and ones developed specifically for deep space communications from the CCSDS and Delay Tolerant Network Research Group (DTNRG).

The OSI model was formalised in 1977 [12], and continues to be deployed in terrestrial networks today. However, while suitable for the terrestrial environment, its operation is not as transparent in interplanetary space. Problems exist with several of the protocols: these limitations are thoroughly discussed in [13] [14].

Stack enhancements have built on the OSI model. Both the CCSDS and DTNRG make structural and protocol modifications, although develop their models in different ways. The contrasting approaches enable comparisons to be made regarding further decisions on protocol stack developments.

4.1. Limitations of the Proposed Interplanetary Stacks

Proposals from the CCSDS and DTNRG are successful standalone protocol stacks. They do, however, possess limitations. The window-based approach to flow and congestion control is a limitation of the CCSDS protocol stack, being difficult to implement successfully where transmission times are long and variable. Indeed, the use of the Transmission Control Protocol (TCP) [15] and modified versions in a protocol stack designed for long-distance operation is not understood, being widely publicised as inadequate to address the challenges in the interplanetary environment [16].

The DTN's bundling protocol may also be classed as a limiting factor. While ensuring reliability, bundling is ineffective when network resources are limited. Store-and-forward techniques are suitable only when transmissions are of lengths suitable for the buffers. The store-and-forward approach is also unsuitable for multimedia applications. While multimedia does not need one hundred percent reliability, it does have strict requirements on a steady traffic rate and bounded jitter. Bundling will not allow a steady transmission rate.

In addition, neither of the CCSDS or DTN protocol stacks contain autonomy or flexibility. While the bundle layer can be deployed on top of various transport protocols, there is no autonomy in protocol selection. In addition, there is no indication in either stacks that the protocols used can change during the course of a transmission.

The limitations of these most recent proposals reveal that research gaps exist in the development of protocols and protocol stacks for the interplanetary environment. Not only are the stacks insufficient for

current missions in space, they are almost certainly insufficient for anticipated missions.

5. A Novel Approach to a Context-Aware Protocol Stack

There are two options to deploying context-awareness. Firstly, a protocol stack may be developed with minimal protocols, each of which is imbued with context-awareness. The protocols can adapt and react to the external environment, the network, and traffic. A second approach involves deploying multiple protocols in each layer, and choosing one in response to application layer knowledge and context information.

The first option is explored in [17], where the case for building a context-aware TCP/IP is evaluated. The context-aware information is provided to the TCP and IP protocols by CATNIP, a Context-Aware Transport/Network Internet Protocol, and the resulting protocols are imbued with decision-making ability. However, while using context knowledge to make decisions, it is context knowledge in a limited form. For example, knowledge includes the number of packets in the last window or the number of packets in a web document transfer. Contexts relating to the external environment are not incorporated into the choice. Using context knowledge in this manner is unlikely to be effective in the interplanetary environment: on sensing context information, the more effective option may be to deploy a UDP-type service. Having a TCP service with context-abilities will therefore be ineffective.

The second approach to protocol selection based on application layer knowledge is deployed by the CCSDS. They propose a protocol suite with three transport layer protocols to respond to varying qualities of service. The protocol called will depend on QoS information passed from the application layer. While the idea of multiple protocols in attempting to satisfy QoS levels is considered effective, the protocols deployed are not. Research reveals that TCP, the protocol used by the CCSDS to provide a service of full reliability, has limitations when deployed in the interplanetary environment. While the CCSDS have extended the traditional TCP protocol, these extensions are insufficient for the diverse situations experienced in deep space.

A protocol suite, shown in Figure 3, has been developed to consider how context-awareness may be incorporated. It exploits the idea of protocol selection based on application layer knowledge and context information, with multiple protocols therefore deployed in each layer.




Figure 3 Remodified Protocol Stack and Suite of Protocols

The proposed stack is a hybrid of stacks from the ISO, DTNRG and CCSDS. In addition to incorporating attractive qualities from each, a layer of context-awareness is included, an element not explicitly featured in other models to date.

The construction of the application layer builds on ideas presented in [18] [19] [20]. In [18], the applications are considered in terms of being delay-sensitive or delay-insensitive. Due to the limited resources in deep space, it is thought necessary to categorise applications further according to their mission-critical importance.

The main revision to this model is the incorporation of a context-aware layer which overlays all decisions made and passes information in the contexts shown in Figure 1. This context-aware layer builds on an idea presented in [20], which includes an Intelligent Internet Protocol (IIP) in the network layer. A context-aware layer has numerous applications in deep space communication. It may be used to reconsider the use of all stack layers depending on the application's needs. As stated earlier, bundling is ineffective where network resources are limited. The context-aware layer could therefore identify if and where network resources are limited and adopt the most appropriate approach in response. The context-aware layer could also set the time-to-live (TTL) field of the IP packet, and it can decide if end-to-end or hop-to-hop flow control is the most appropriate approach. Considering the decisions required in interplanetary communication, the incorporation of a context-aware layer is overdue.

Information from the application layer regarding QoS requirements is used along with context information to make transport, security, network, data link and physical layer choices. These layers are populated with various protocols in the attempt of ensuring that the most appropriate protocols are selected based on the additional information, an approach modelled on the CCSDS protocol stack. The context-aware layer represents an addition to all protocol stacks. The idea of incorporating context-awareness is to empower the stack to make communication choices based on network resources, application traffic, and environmental characteristics, as discussed earlier in Figure 1.

The inclusion of the security layer results from references made to [19]. Due to the anticipated increase in mission frequency and importance, security is needed to ensure the authenticity, integrity and confidentiality of all transmissions. Its incorporation becomes even more important with the intended increase in manned missions.

The DTN proposal introduces several concepts which are included in the remodified stack. These include the convergence layer, the bundling protocol and the Licklider Transport Protocol (LTP) [21]. The idea of a convergence layer is attractive, with its ability to hide characteristics of networks built using different principles. The store-and-forward approach to message delivery has benefits for ensuring transmission for applications with reliability requirements. The LTP is incorporated as a result of its non-reliance on IP. Durst et al identify that many

specialised networks cannot and will not ever run IP [22]. Therefore, the incorporation of this factor into the model was required.

An understanding of transmission capabilities of the physical layer is required. Regardless of performance improvements in the stack, the physical layer will always limit the speed of communication. Transmission rates envisaged for the 2009 launch of the Mars Telecommunications Orbiter are discussed in [23]: radio waves will reach ultra high frequency (UHF) speeds of 128 kbps and microwaves within the X- and Ka-bands will reach 4 Mbps [23]. While optical communication is desired to achieve ultimate communication speeds, the speed of light is a limitation in long-distance communication. Therefore, regardless of performance improvements achieved by refining the protocol stack, communication over long-distances will never achieve true real-time responses.

Paths through the protocol stack in Figure 3 are indicated for missions with different quality of service requirements. The QoS required for the transmission of scientific data is low: there are no strict latency requirements, resulting in the ability to deploy all flow and congestion control procedures.

Tracking the path through the protocol stack, the bundling protocol and associated Licklider Transport Protocol will be selected. This will enable flow and congestion control on a hop-by-hop basis. The transmission can apply security: SCPS-SP [24] is selected. As LTP removes the need for a network protocol, the data link layer is encountered next. TM Space is selected, being the link layer protocol developed for transmitting telemetry from spacecraft to ground. Finally, the transmission is pushed from the protocol stack on a radio wave.

At the other end of the QoS spectrum, the transmission of telemetry is considered. This application has strict boundaries on transmission latencies. In addition, it has a low tolerance of data loss. Therefore, its path through the protocol stack will be different. On leaving the application layer, the bundling layer is bypassed and TP-Planet [16] is selected from the transport layer, with the assumption that the end-to-end approach is more appropriate in this case to minimise additional processing time. Due to the need to minimise additional processing, the security layer is bypassed. The network layer protocol selected is SCPS-NP [24], and the data link protocol selected is TC Space [24]. Finally, the signal is pushed from the stack on a radio wave link.

The efficiency of operating a protocol stack in this manner cannot be guaranteed. Initial insights may indicate that any benefits from the incorporation of context-awareness will be removed by the additional processing required through the selection of the most

appropriate protocols. Exploratory simulations will evaluate the proposal.

The model in Figure 3 is consequential of ideas on flexibility requirements in interplanetary communication, awareness of interplanetary protocol capabilities, and several assumptions. Assumptions provision for worst case scenarios; these include the absence of end-to-end connectivity, excessive latencies, high error rates, asymmetric data rates and lack of a fixed infrastructure. In deep space, these occurrences currently have a high probability of affecting communication. While latency cannot be overcome, it is anticipated that the effects of the other factors will be minimised as infrastructure deployments become more common. With regard to latency, protocol developments are working to mask its effects where possible.

The protocol stack is also a consequence of the available communication protocols. Research has occurred on the development of interplanetary transport protocols, although each researcher defines the optimal protocol differently. We (the authors) have similar unique ideas, and feel that, to incorporate reliability into the protocol suite, it is necessary to use functions from several protocols.

TP-Planet is promising as a reliable rate-based transport protocol. Its approach to connection establishment is effective; it sends the initial data transmission with the connection request, maximising throughput while ensuring reliability. Its Blackout State algorithm is also attractive, called when the communication encounters periods of intermittent connectivity. However, TP-Planet is a 'heavy' protocol, having a thorough set of flow and congestion control procedures. In certain circumstances, these protocol features must be dropped; where bandwidth is limited, it is inappropriate to perform retransmissions. A more lightweight protocol will be required with minimal requirements to ensure that, in the best-case scenario, that the application QoS requirements are met and in the worst-case, that the data reaches its destination.

Therefore, while Figure 3 shows multiple protocols in each layer, it is anticipated that a future protocol stack will deploy single protocols in each layer with functionalities sufficient to cope with all interplanetary eventualities. Attractive features from several protocols have been identified, which could be combined into a single protocol. However, deploying all attractive features would result in conflicting operations. TP-Planet, for example, uses congestion control while SCPS-TP (UDP Option) does not. However, it is intended that functions will be selected in response to application layer knowledge, context awareness and time. Figure 4 illustrates how the protocols may be used.




Figure 4 Protocol Functions with Distance from Source

Figure 4 represents a protocol stack encountered at each hop in an interplanetary transmission. It is intended that a single transport protocol is deployed, with elements from several different protocols. Protocol functions can be called in response to external information. If the transmission cannot cope with loss, is delay-insensitive, bandwidth is unlimited and connectivity continuous, a reliable transport protocol can be deployed. Flow and congestion control from TP-Planet can be called, when the instruments network via IP. TCP-Peachtree's [25] multicasting capability can be called if required. With time (and distance), the network conditions are likely to change, and the ability to transfer data and achieve QoS will become more restricted. For this reason, at the next hop, a transport protocol with less functionality will be called. In Figure 4, MDP is chosen. The ability of Freeze-TCP [26] to handle long and frequent disconnections may also be included as a precaution. The effect of calling MDP is to continue to implement a reliable service but to minimise excess data. If, by the next hop, the signal has not reached its destination, a protocol with even less functionality may be used. In Figure 4, SCPS-TP (UDP Option) is called, a protocol which does not implement flow or congestion control. Similar protocol selections may occur in the remaining protocol layers. The model also considers the case when the network protocol is not IP: in this case,

bundling and LTP may be called. Where the transmission is multimedia, RCP-Planet [27] may be used.

Adaptive protocols are thought necessary, given the dynamic nature of the interplanetary environment. Due to the vast distances over which a signal may travel, changing network conditions must be expected. Although using such a protocol will be time-consuming, it is anticipated that any costs will be outweighed by the benefits.

6. Conclusions and Future Work

Certain properties of the interplanetary environment seem insurmountable [28]. This is true, to a degree. Regardless of any stack and protocol developments, propagation delay in the interplanetary environment cannot be overcome. Deep space is an environment in which the speed of light becomes a limiting factor. Other factors, which seem too large to control, can in fact be overcome. The occurrence and effects of intermittent connectivity, high error rates, asymmetric data rates and lack of a fixed infrastructure can be minimised by the deployment of additional components and more carefully designed protocols and protocol stacks. This is a gap which Cisco can fill, although the race has now begun.

Our work involves developing a transport protocol for operation over long-distances, to be implemented within a context-aware protocol stack. Research has identified attractive features from currently proposed protocols, and a gap exists with regard to extracting the qualities from each and combining them together.

Future work will continue with the development of a simulation environment during a summer internship with Cisco Systems. Simulation scenarios will evaluate protocol features under a variety of network conditions over long-distances. The extension of MPLS into space is a key area to be investigated, having been highly successful in enabling next generation networking on Earth. However, making the protocol more lightweight is required before its deployment can be considered. Simulations will indicate the features to include in the redesigned protocol and the applicability of the proposed stack. However, measuring the effectiveness of these procedures is only one stage in considering the components of the optimal protocol. The risk that the selected features will have interoperability issues is also anticipated.

References (URLs correct as of April 2005)

- [1] National Aeronautics and Space Administration, *"The New Age of Exploration: NASA's Direction for 2005 and Beyond"*, February 2005.
- [2] Hogie, K; Criscuolo, E; Parise, R, *"Using Standard Internet Protocols and Applications in Space"*, Computer Networks 47, 2005, Pages 603 – 650.
- [3] CANDOS, <http://ipinspace.gsfc.nasa.gov/CANDOS>
- [4] Postel, J, *"User Datagram Protocol"*, Request for Comments 768, August 1980.
- [5] Macker, J, P; Adamson, R, B (1999) *"The Multicast Dissemination Protocol (MDP) Toolkit"*, IEEE, Pages 626 – 630.
- [6] Jackson, J, *"The Interplanetary Internet"*, IEEE Spectrum, August 2005, Pages 30 – 35.
- [7] Floreani, D, Wood, L, *"Internet to Orbit"*, Cisco Systems, Third Quarter 2005, Pages 19 – 23.
- [8] Consultative Committee of Space and Data Systems, www.ccsds.org
- [9] Khedr, M, Karmouch, A, *"ACAN – Ad Hoc Context Aware Network"*, IEEE 2002 Canadian Conference on Electrical and Computer Engineering, Volume 3, 12 – 15 May 2002, Pages 1342 – 1346.
- [10] Curtis, S, A, Rilee, M, L, Truszkowski, W, Clark, P, E, *"ANTS for the Human Exploration and Development of Space"*, Proceedings of the IEEE 2003 Aerospace Conference, Volume 1, March 8 – 15, 2003, Pages 1 – 7.
- [11] D'Arrigo, P, Santandrea, S, *"APIES: A Mission for the Exploration of the Main Asteroid Belt Using a Swarm of Microsatellites"*, Acta Astronautica, 2005, In Press, Corrected Proof.
- [12] Tillman, M, A, Yen, C-C, *"SNA and OSI: Three Strategies for Interconnection"*, Communications of the

ACM, February 1990, Volume 33, Number 2, Pages 214 – 224.

- [13] Akyildiz, I, F, Akan, O, B, Chen, C, Fang, J, Su, W, *"Interplanetary Internet: State-of-the-Art and Research Challenges"*, Computer Networks 43, 2003, Pages 75 – 112.
- [14] Hogie, K, Criscuolo, E, Parise, R, *"Link and Routing Issues for Internet Protocols in Space"*, IEEE Proceedings of the 2001 Aerospace Conference, Volume 2, 10 – 17 March 2001, Pages 2/963 – 2/976.
- [15] Postel, J, *"Transmission Control Protocol"*, Request for Comments 793, September 1981
- [16] Özgür, B, A, Fang, J, Akyildiz, I, F, *"TP-Planet: A Reliable Transport Protocol for Interplanetary Internet"*, IEEE, 2002.
- [17] Williamson, C, Wu, Q, *"A Case for Context-Aware TCP/IP"*, ACM Sigmetrics Performance Evaluation Review, Volume 29, Issue 4, March 2002, Pages 11 – 23.
- [18] Hogie, K; Criscuolo, E; Parise, R, *"Using Standard Internet Protocols and Applications in Space"*, Computer Networks 47, 2005, Pages 603 – 650.
- [19] Burleigh, S, Cerf, V, Durst, R, Fall, K, Hooke, A, Scott, K, Weiss, H, *"The Interplanetary Internet: A Communications Infrastructure for Mars Exploration"*, 53rd International Astronautical Congress, 10 – 19 October 2002.
- [20] Malakooti, B, Robinson, B, Thomas, I, Bhasin, K, Holtz, A, *"An Intelligent Internet Protocol Architecture for Space-Based Networks"*, Case Western Reserve University, Department of Electrical Engineering and Computer Science, NASA Glenn Research Centre, 2004.
- [21] Ramadas, M; Burleigh, S; Farrell, S; *"Licklider Transmission Protocol – Specification"*, Delay Tolerant Networking Research Group Internet Draft, Ohio University, NASA/Jet Propulsion Laboratory, Trinity College Dublin, July 2005.
URL: <http://www.ietf.org/internet-drafts/draft-irtf-dtnrg-ltp-03.txt>
- [22] Durst et al, *"The DTNRG: Where Are We Now?"*, The MITRE Corporation, Intel Berkely Research Lab, 2005.
- [23] Edwards, C, Gibbs, R, *"Focused Discussions – Missions in Definition: Mars Telecommunications Orbiter 2009"*, National Aeronautics and Space Administration, 2006.
- [24] Consultative Committee for Space Data Systems, *"Overview of Space Link Protocols"*, Green Book, June 2001.
- [25] Akyildiz, I, F; Fang, J, *"TCP-Peachtree: A Multicast Transport Protocol for Satellite IP Networks"*, IEEE Journal on Selected Areas in Communications, Volume 22, Number 2, February, 2004, Pages 388 – 400.
- [26] Goff, T; Moronski, J; Phatak, D, S; Gupta, V *"Freeze-TCP: A True End-to-End TCP Enhancement Mechanism for Mobile Environments"*, IEEE Proceedings of the Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies, 26 – 30 March, 2000, Volume 3, Pages 1537 – 1545.
- [27] Akyildiz, I, F, Fang, J, *"RCP-Planet: A Rate Control Scheme for Multimedia Traffic in Interplanetary Internet"*, Broadband and Wireless Networking Laboratory, 2003.
- [28] Fall, K, *"A Delay-Tolerant Network Architecture for Challenged Internets"*, Intel Research Berkeley, February 2003.