Extending Internet into Space – ESA DTN Testbed Implementation and Evaluation

Christos Samaras[1], Ioannis Komnios[1], Sotirios Diamantopoulos[1], Efthymios Koutsogiannis[1], Vassilis Tsaoussidis[1,*] Giorgos Papastergiou[2], and Nestor Peccia[3]

[1] Democritus University of Thrace
Department of Electrical and Computer Engineering,
12 Vas. Sofias Str., 67100 Xanthi, Greece
{csamaras,ikomnios,sdiaman,ekoutsog,vtsaousi}@ee.duth.gr
[2] Hellenic Aerospace Industry S.A.
P.O. Box 23, 32009 Schimatari, Greece
{papastergiou.georgios@haicorp.com}
[3] ESA/ESOC
Robert Bosch Strasse 5, D-64293 Darmstadt, Germany
{nestor.peccia@esa.int}

Abstract. As the number and complexity of space missions increases, space communications enter a new era, where internetworking gradually replaces or assists traditional telecommunication protocols. The Delay Tolerant Network (DTN) architecture has recently emerged as a communication system for challenged networks, originally designed for the Interplanetary Internet. In the context of our project with ESA called "Extending Internet into Space - ESA DTN Testbed Implementation and Evaluation" we intend to deploy a distributed, flexible and scalable DTN testbed for space communications. The testbed will provide the supportive infrastructure for the design and evaluation of space-suitable DTN protocols, architectures, and routing policies to allow efficient deep-space communications. Throughout the project, we will demonstrate the operational capabilities of the DTN protocols in space; design and evaluate novel transport protocols and architectures for reliable data transfer in space; and investigate routing algorithms that comply with ESA's policies and resource status.

Keywords: Delay-Tolerant Networking, Testbed, Deep-Space Communications.

1 Introduction

Currently, all space communications are static, inflexible, and involve prior scheduling of communication contacts. In this context, less sophistication was required from communication protocols; the link layer was the dominant layer for space communications; routing was never an issue; end-to-end reliability was frequently overlapping

* Technical/Scientific Leader.

F. Granelli et al. (Eds.): MOBILIGHT 2009, LNICST 13, pp. 397–404, 2009.

with reliability of a single hop; congestion and overflow were absent due to strict scheduling of communication activities and admission control; and the limited required sophistication was shifted to the application layer.

However, future missions become more complex and new communication architectures and protocols for backbone, access, and proximity networking need to be designed, validated and optimized. Two new major properties have changed the spectrum of potential architectural choices for space communications: (i) the multi-hop architecture, which is required to reach deep space and (ii) the increasing number of alternative communication paths that may be used to reach a single receiver. Along these two properties, the demand for interoperability among space agencies has also contributed towards the emerging field of Delay-Tolerant Networking (DTN) [1]. DTN architecture is essentially a communication system to provide data transfer services in challenged environments, featuring extreme operational characteristics such as high propagation delays or network partitions. DTN applicability spreads over a wide spectrum of networking environments. With respect to space, DTN is envisioned to support Internet-like services across interplanetary distances.

As DTN becomes a standard architecture included in the Consultative Committee for Space Data Systems (CCSDS) standardization procedures, a testing and verification infrastructure emerges along with a set of scenarios, operations and evaluation procedures. In this context, we are designing and building an appropriate DTN testbed to evaluate associated scenarios, mainly targeting Mars-to-Earth communications (for an example scenario, see Fig. 1). This testbed will allow for cost-effective evaluation and optimization of space communication designs for reliable and efficient data delivery, expose the corresponding constraints, and uncover potential tradeoffs. The testbed will verify the robustness of space communications network, which can potentially reduce mission risks and enhance space data transfer rates.

The remainder of the paper is organized as in the following. In Section 2, we elaborate on current space communications, which are mainly supported by CCSDS protocols, and discuss the DTN architecture. Our DTN testbed is presented in Section 3. We outline the goals and research directions of our project in Section 4. Finally, we conclude the paper in Section 5.

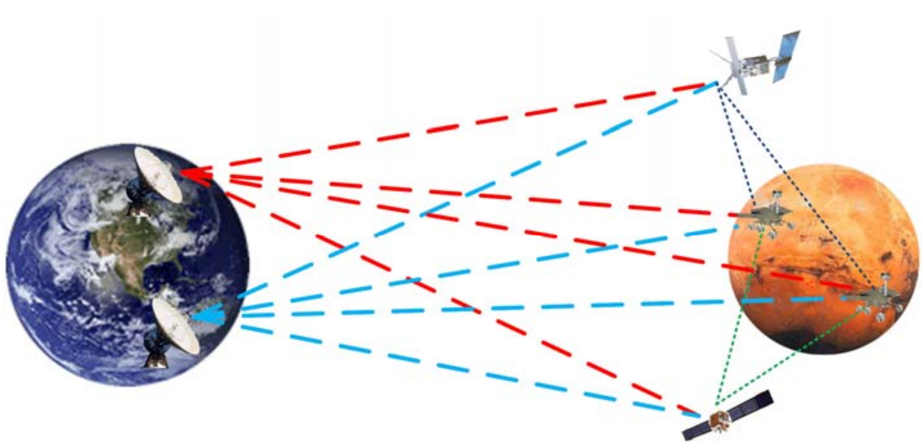

Fig. 1. Space Network Environment

2 Space Communications and DTN Architecture

2.1 Space Communications Today

Existing space communication infrastructure is built and managed around the operational targets of space missions and allows mainly for sending commands from mission control centers to spacecraft, and for receiving telemetry data from spacecraft to ground. In this context, the link layer is today the dominant layer for space communications. CCSDS has established a wide range of standards for space communications, including Telecommand and Telemetry Space Data Link Protocols [2]-[3], and Space Packet Protocol [4]. No provision is made for direct communication among space elements which would enable space internetworking using different layers at different points in end-to-end data flow.

Some technological advances have already been made towards the direction of end-to-end services. A recent achievement is, for example, the CCSDS File Delivery Protocol (CFDP) [5], which is designed for reliable file transfer across interplanetary links. CCSDS has also enabled onboard systems to have their own IP addresses. This is accomplished by either direct use of IP or an abbreviated form of IP that is the Network Protocol (NP) component of a four-layer stack of protocols, known as the Space Communication Protocol Standards (SCPS) [6]. Both of these capabilities allow dynamic routing through different paths in a connectionless fashion. However, it is inherent in the physics of space communications that both SCPS-NP and IP can not work well in areas of disrupted and long-delay communications that mainly characterize deep-space communications; IP is mainly bottlenecked by the conflicting nature of dynamic routing that requires frequent information exchange and the prohibiting nature of space communications that confine the concept of synchronization and accuracy.

Although the main target of CCSDS is to create a common protocol stack to be used by all agencies in the missions they plan, this is not what happens today, as the communication protocols used are mainly mission-specific. Inter-agency communication requires a common interoperable platform to allow for more natural communication that may replace a series of encapsulation and tunneling patches and help realize joint missions or multiple parallel missions of different space agencies. Furthermore, resource sharing and store-and-forward practices allow for enhanced connectivity and reduced mission costs. Typically, interoperability requires a level of convergence where different protocols can be deployed above or below that level. The experience from the Internet shows that interoperability is an issue along with cost and communication efficiency.

Another problem, present in space missions today, is the manual operation of communications to and from spacecraft. Indeed, no automation has been imported into the way space data transfers take place. Hence, as more and more assets are sent into space, scheduling of transmissions becomes too complex to be manually operated. That said, the need for deployment of an internetworking architecture in space is becoming more urgent than ever.

2.2 Delay-Tolerant Networking as a Candidate Architecture for Space

Delay/Disruption-Tolerant Networking has been proposed to overcome relevant problems that arise from current deep-space communication networks, such as absence of a common layer for interoperability among all space agencies; inefficiency in exploiting all communication opportunities; limited number of alternative communication paths that may be used to reach a single receiver; and the static and human-operated mission control. The architecture of DTN is based on, but not limited to that of the Interplanetary Internet [7]-[8]. DTN architecture embraces the concept of occasionally-connected networks that may suffer from frequent partitions and that may be comprised of more than one divergent set of protocols or protocol families. Nowadays, DTN has emerged as a recognized networking research area and has become an interesting idea for challenging environments within the terrestrial Internet as well. For example, DTNs are expected to provide connectivity to the edges of the current Internet infrastructure. DTN, as overlay architecture, provides Internet like services and is capable of extending store-and-forward architecture with permanent storage capabilities. Deep space missions are expected to sufficiently exploit an infrastructure that allows the efficient communication between in-space entities, such as explorer spacecraft, landed vehicles, and orbiters.

DTN employs an end-to-end message-oriented overlay, the so-called bundle layer, that exists at a layer above of the transport (or other) layer of the networks on which it is hosted and below applications. Devices implementing the bundle layer are called DTN nodes. A DTN-enabled application sends messages of arbitrary length that are transformed by the bundle layer into one or more protocol data units, called bundles. Bundles are thereupon forwarded by DTN nodes towards their destination.

Persistent storage (such as disk, flash memory, etc.) is employed to address discontinuous end-to-end connectivity, as it does not expect that network links are always available or reliable. The DTN architecture provides two features for enhancing delivery reliability: end-to-end acknowledgements and custody transfer. The latter mechanism permits delegating the responsibility for reliable transfer among different nodes in the network. When custody transfer is requested, the bundle layer employs a coarse-grained timeout and retransmission mechanism and an accompanying custodian-to-custodian acknowledgement signaling mechanism.

Therefore, custody transfer allows the source to assign retransmission responsibility and recover its retransmission-related resources relatively soon after sending a bundle. Not all nodes are required by the DTN architecture to accept custody transfers. Furthermore, a DTN node may have sufficient storage resources to sometimes act as a custodian, but may elect not to offer such services when congested or running low on power.

3 Testbed Architectural Design

Our testbed design emphasizes mainly on the properties of communication within deep space. We refer to our system as *DTN testbed* or simply *testbed*. During the initial deployment of the DTN testbed, our system will include modeled network links and real protocol implementations that operate at the network layer and above. Thus, no specialized hardware will be needed, all network parameters will be easily controlled, and the results will be reproducible. This will allow for initial cost-effective protocol evaluation and validation, and for smoother testbed extensions discussed later in the paper.

3.1 DTN Testbed Design Goals

The main design goals of the DTN testbed regarding its accuracy and efficiency are:

(i) *Dynamic control of network parameters*. The testbed should be able to emulate fundamental network parameters (such as bandwidth, packet error rate, propagation delay, and available connectivity), and adapt realistically and dynamically to changes in those parameters in real-time.

(ii) *Scalability*. Although current deep-space communications involve a limited number of communication nodes, the testbed should be able to scale well over a larger number of communication nodes to allow for emulation of any future deep-space communication scenarios, which will include several planetary surface networks and relay satellites.

(iii) *Transparency*. Network emulation should be transparent to upper layer protocols and applications, thus permitting their use and evaluation without need for modification.

(iv) *Flexibility*. The testbed should be flexible enough to: emulate any space communication topology; incorporate new protocols, applications and mechanisms; interoperate with other similar DTN testbeds; and provide a reusable infrastructure towards an actual hardware testbed.

3.2 Architecture

In this section we describe the functionality of basic testbed elements and analyze the interaction among them. In Fig. 2, we present the basic structure of the proposed testbed and the interconnection between testbed components.

The *Scenarios Description and Results Visualization System* consists mainly of two components: a tool for describing deep-space communication scenarios, and a tool for

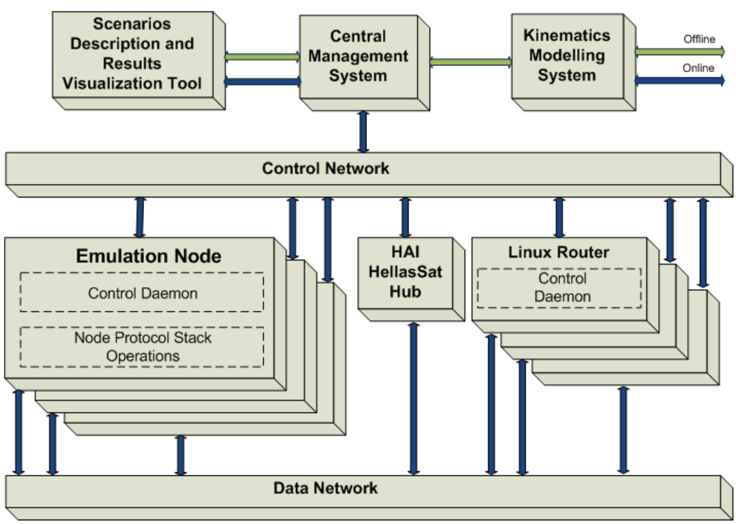

Fig. 2. DTN Testbed Architecture

visualizing emulation results. A *scenario description* contains information about the nodes in the network (e.g., landers, rovers, relay satellites, ground stations, and Mission Operation Centers), including initial position and movement (when applicable), communication protocols that will be used, link characteristics (e.g., frequency band, bandwidth, transmission power, etc.) and alike emulation parameters. The Results Visualization tool is used to present performance results graphically either in real-time (online) or after the emulation ends (offline), collected by all testbed elements through the control network. Performance results include file delivery time, retransmission overhead, status reports, etc. The Scenario Description tool passes scenario information to the Central Management System, while the latter forwards collected results to the Results Visualization tool.

The emulation system is divided into two discrete networks, namely a *Control Network* and a *Data Network*. The former is responsible for coordination of the emulation via control messages, while the latter moves data (such as output data) within the testbed. *Emulation nodes* represent distinct network nodes in the system, and a number of routers can be exploited to emulate multiple links, as is the case of a planetary surface network for example. The final component of the testbed is a real geostationary link through *Hellas Sat GEO Satellite* that will be integrated into the DTN testbed for testing purposes.

4 Goals and Research Directions

Core target of our DTN testbed is validation of the DTN architecture in space. Indeed, our project will serve as a means to reveal problems and deficiencies of current DTN specification [9], and propose relevant adaptations. DTN architecture will be tested in terms of: applicability to space networks; conformance with current space communication protocols and ESA's infrastructure; and performance.

Since manual routing as currently operated in space is costly and not scalable, more flexible routing schemes need to be established. Common IP-based routing protocols cannot operate in space, given that they rely upon constant network connectivity. A DTN-compatible routing algorithm should make communication decisions on the basis of locally available information such as: communication opportunities; expiration-time of messages to be delivered; performance history of communication paths, and analogous heuristics. Thus, static routes can be defined beforehand, but dynamic routes can also be discovered (for example, see [10]).

One major scope of our DTN testbed is to implement and evaluate protocols and mechanisms that enhance interoperability among space agencies, while also preserving their individual policies. In that context, a routing scheme based on priorities will be implemented. Indeed, in space, data priority can be associated with route priority. For instance, an agency may prefer to wait for communication links between its own assets, instead of forwarding data via some other agency's space systems, even if the latter contact opportunity becomes available sooner. Alternatively, an agency may decide that connectivity through its own resources may delay considerably, and go for a shortest path using another agency's resources. Such decision may fit to cases where high-priority data should be promptly forwarded. Inter-agency agreements can lay the ground for resource sharing between space missions, and DTN routing policies can exploit the offered contact opportunities.

The Delay-Tolerant Networking architecture essentially relies upon underlying network services adapted to the special networking conditions. In the context of deep-space communications, novel transport and application layer algorithms should be established for efficient and reliable data transfers. We have implemented and evaluated through ns-2 simulations [11] two protocols to address reliable data transports in space: Deep-Space Transport Protocol [12] and Delay-Tolerant Transport Protocol [13]. Deep-Space Transport Protocol (DS-TP) utilizes the hop-by-hop, store-and-forward message switching principle that governs today's space communications, and mitigates the need for congestion avoidance. DS-TP's novel, proactive retransmission scheduling allows for efficient and fast retransmission of corrupted packets due to high BERs or blackouts. Delay-Tolerant Transport Protocol (DTTP), like DS-TP, provides reliable data transfer over challenged network environments. It comprises a packet-oriented transfer approach over multi-hop and collaborative end-to-end paths, and acquires available bandwidth resources via its rate-based transmission behavior. DTTP can either operate as a standalone transport protocol (facing long delays and disruptions) or complement DTN architecture in space.

In the context of our project, we plan to further enhance DTTP and DS-TP functionality. Capabilities to be added include (but are not limited to) the following:

(i) *Proactive retransmission*. We will implement and evaluate retransmission mechanisms, that inject redundant data into the network to strengthen reliable transfers against packet losses. Alternatively, we will deploy packet-level erasure coding to allow for advanced recovering from packet losses. Coding rate can be dynamically adjusted, based on network measurements. We will investigate the trade-off between processing overhead, retransmission overhead, and recovering capability.

(ii) *Parallel data transfer*. Data transfer can be accomplished in parallel data paths, exploiting various communication opportunities. Sequence of application data is resumed at the receiver. Parallel data transfer can be implemented by preserving the original sequence number space. This feature requires explicit definition in the protocol header, so that the final destination can anticipate and merge data packets coming from different paths.

(iii) *Sending rate adaptivity*. In space settings, sending rate can be accurately characterized as temporarily constant. Sending rate adaptation can follow network events such as storage capacity exhaustion. These network events can be either perceived by senders through advanced mechanisms or explicitly signaled by receivers.

5 Conclusions

Currently, all space communications are static, inflexible, and manually preconfigured long before they actually take place. Network disconnections, potentially huge propagation delays, high link error-rates, and bandwidth asymmetries compose the space networking environment. Delay-Tolerant Networking is a candidate communication architecture in space, as it can stand long delays and network partitions. In this context, we construct a DTN testbed that integrates a variety of tools and protocols in order to emulate realistic deep-space communications scenarios with varying network elements, topologies, and operational parameters. The testbed will form a valuable platform for testing DTN architecture in space, and a research framework for evaluating new mechanisms and protocols.

References

1. Cerf, V., et al.: Delay-Tolerant Network Architecture, IETF RFC 4838, information (April 2007), http://www.ietf.org/rfc/rfc4838.txt
2. TC Space Data Link Protocol. Recommendation for Space Data System Standards. CCSDS 232.0-B-1. Blue Book. Issue 1. Washington, D.C. CCSDS (September 2003)
3. TM Space Data Link Protocol. Recommendation for Space Data System Standards. CCSDS 132.0-B-1. Blue Book. Issue 1. Washington, D.C. CCSDS (September 2003)
4. Space Packet Protocol. Recommendation for Space Data System Standards. CCSDS 133.0-B-1. Blue Book. Issue 1. Washington, D.C. CCSDS (September 2003)
5. CCSDS File Delivery Protocol. Recommendation for Space Data System Standards. CCSDS 727.0-B-2. Blue Book. Issue 2. Washington, D.C. CCSDS (October 2002)
6. Space Communications Protocol Specifications (SCPS), http://www.scps.org/
7. Burleigh, S., Hooke, A., Torgerson, L., Fall, K., Cerf, V., Durst, B., Scott, K.: Delay-Tolerant networking: an approach to InterPlaNetary Internet. IEEE Communications Magazine 41(6), 128–136 (2003)
8. Cerf, V., et al.: Delay-Tolerant Network Architecture: The Evolving Interplanetary Internet, IPNRG Internet Draft draft-irtf-ipnrg-arch-01.txt
9. Scott, K., Burleigh, S.: Bundle Protocol Specification, IETF RFC 5050, experimental (November 2007)
10. Burleigh, S.: Dynamic Routing for Delay-Tolerant Networking in Space, Flight Operations SpaceOps 2008, Conference (Hosted and organized by ESA and EUMETSAT in association with AIAA) (2008)
11. ns-2. The Network Simulator, http://www.isi.edu/nsnam/ns/
12. Psaras, I., Papastergiou, G., Tsaoussidis, V., Peccia, N.: DS-TP: Deep-Space Transport Protocol. In: IEEE Aerospace Conference 2008, Montana, USA (March 2008)
13. Samaras, C.V., Tsaoussidis, V.: DTTP: A Delay-Tolerant Transport Protocol for Space Internetworks. In: 2nd ERCIM Workshop on eMobility, Tampere, Finland (May 2008)