Evaluation of CCSDS File Delivery Protocol over Delay Tolerant Networks

Efthymios Koutsogiannis
Department of Electrical and Computer Engineering
Democritus University of Thrace
Xanthi, GREECE
Email: ekoutsog@ee.duth.gr

Giorgos Papastergiou
Hellenic Aerospace Industry S.A.
Schimatari, GREECE
Email: papastergiou.georgios@haircorp.com

Vassilis Tsaoussidis
Department of Electrical and Computer Engineering
Democritus University of Thrace
Xanthi, GREECE
Email: vtsaousi@ee.duth.gr

Abstract— **Space communications enter a new era where a multihop architecture is being exploited and an increasing number of alternative communication paths may be available. Current space applications, such as CFDP for file transfer, rely on static routing and can not efficiently perform in these challenging environments. The most promising solution is the emerging Delay Tolerant Networking (DTN) Architecture and the accompanying Bundle Protocol which allow for reliable store-and-forward message switching, custody transfer and dynamic routing. However, a space-oriented file transfer protocol for DTN is required. In this work, we integrate CFDP with ION (a Bundle Protocol implementation produced by the Jet Propulsion Laboratory) and evaluate its performance in a DTN testbed.**

Keywords— **DTN, CFDP, ION, testbed, evaluation**

I. Introduction

The increasing interest in exploration of Space over the last years has led to the launch of several communicating devices. Therefore, automated communication with and among space equipment constitutes a necessity. Mechanisms and protocols that support routing and transfer reliability under conditions of intermittent connectivity and long propagation delays need to be used. These new mechanisms must also cooperate with existing – and in many cases obsolete – infrastructure.

Even though there has been much research on space communication protocols, most protocols appear to have a rather cross-layer functionality. Some protocols (e.g., Space Packet Protocol [1]) appear to have naming functionalities supporting static routing and assume that reliability is granted due to link layer protocol, while others include transport layer functionalities as well (e.g., CCSDS File Delivery Protocol [2]). Space agencies have not yet agreed upon a common architecture, since each agency uses mainly its own equipment; therefore, protocols were designed with very strict specifications and for agency-exclusive use. The distinction between functionalities at each layer will lead to a structured and coherent protocol design. Recently, space agencies agreed to cooperate and share resources in the future; this calls necessarily for a common architecture. Delay Tolerant Networking specifically designed for space conditions is widely adopted and tends to dominate the field.

From an engineering point of view, designing new protocols from scratch based on previous networking experience may be the most prominent solution, but from a financial point of view this is prohibitive for space agencies. The design and testing of space communication protocols requires many years of effort and huge amount of money to be spent. For this reason, space agencies prefer to adjust and enhance current operational protocols, rather than to implement new protocols from scratch. This approach may impose functionality overlapping among protocol's components and any redundant mechanism should be disabled. Space protocols are typically manually operated, so turning off some of their functionality is quite feasible.

In this paper, we investigate the performance of a file transfer protocol over an implementation of Delay Tolerant Network (DTN) architecture. CCSDS File Delivery Protocol [2], is designed for deep space communication, providing transfer reliability itself, but lacking dynamic routing capabilities. On the contrary, DTN architecture [3] and the accompanying Bundle Protocol [4] is designed to perform efficiently, in the field of dynamic routing in the challenging environment of Deep Space, providing transfer reliability as well. The evaluation of the combination of CFDP as a file protocol and DTN architecture for routing is our main purpose.

II. Protocol Description

In this section we present briefly the main functionalities of CFDP, a protocol designed to meet the needs of space missions for file transfers to and from an onboard mass memory, and Interplanetary Overlay Network (ION, [5]), a space-oriented implementation of the Bundle Protocol, including convergence layer protocols and adapters.

A. CFDP

CCSDS File Delivery Protocol is an Application Layer protocol, but it also provides Transport Layer functionalities, such as detecting and retransmitting corrupted or lost data. Therefore, it can run over an unreliable network protocol such as, the Space Packet Protocol [1] and UDP.

File transmission can be executed reliably (acknowledged mode) or unreliably (unacknowledged mode). CFDP provides file delivery services (i) across a single link (referred to as Core Functionalities) and (ii) over more complex topologies, where CFDP provides subsequent transmissions of files between intermediate nodes, which end up to the destination node (i.e., Extended Procedures / Store-and-Forward Overlay). CFDP includes four modes for sending Negative Acknowledgments (i.e., Deferred, Immediate, Prompted and Asynchronous) and uses positive Acknowledgments (ACKs) as well, to ensure the receipt of critical PDUs.

Even though CFDP has been designed to deal with the space-ground communication constraints, it appears to be unsuitable when communication complexity increases (e.g., dynamic routing or file distribution through multiple relay points in parallel). Therefore, a solution would be to integrate CFDP with the Bundle Protocol, and encapsulate each CFDP PDU in a 'bundle' (a bundle protocol datagram). Thus CFDP PDUs are independently routed in store-and-forward fashion through the network. In this case the CFDP is active only at the endpoint entities and the function of point-to-point retransmission devolves to end-to-end acknowledgement.

B. ION

ION is produced by the Jet Propulsion Laboratory and partially maintained by Ohio University. ION is an implementation of (i) the Bundle Protocol, which is used as an overlay network protocol, implemented as described in Internet RFC 5050 [4], (ii) the Asynchronous Message Service (AMS), and iii) the Licklider Transmission Protocol (LTP) [6] which provides retransmission-based reliability.

The main features of the Bundle Protocol are: (i) store-and-forward message switching, (ii) permanent in-network storage, (iii) custody transfer, (iv) naming and late binding, (v) routing and (vi) reactive and proactive fragmentation.

ION supports both UDP and LTP as bundling convergence layers, as well as two TCP convergence layers (simplified-TCP and the Bundle Relay Service). Additionally, ION supports Contact Graph Routing (CGR, [7]), which enables dynamic routing, where dynamic routes are computed through a time-varying topology of scheduled communication contact. A reference Bundle Protocol implementation, namely DTN2, also exists and is supported by the Delay Tolerant Networking Research Group (DTNRG, [8]). Currently, only the UDP convergence layer of DTN2 interoperates with ION.

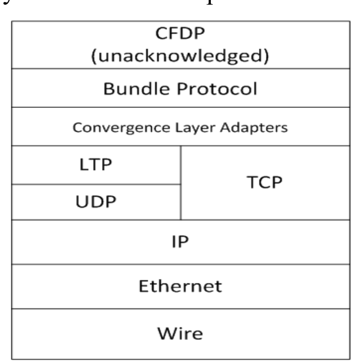

Figure 1. Testbed's protocol stack

III. DISCUSSION

In this section we focus on the transport layer functionalities and we initially analyze the advantages and caveats of CFDP and BP/LTP for data transferring over deep-space links independently, discuss their (re)transmission mechanisms and present any functional similarities among them. Then, we discuss the most important factors that affect the integration of CFDP into ION and the trade-offs among them.

In high error-prone environments, where large chunks of data have to be retransmitted with long and variable propagation delays such as deep-space environments, protocol's (re)transmission tactics play an important role in its performance. Indisputably, retransmission timers are critical components for efficient operation. CFDP does not include any sophisticated mechanism for retransmission timers' calculation. For every critical PDU transmission, such as EOF, Finish and NAK PDUs a preconfigured fixed timer is set. This may lead to late or unnecessary retransmissions and limits protocol's flexibility. Unlike CFDP, LTP (at least as it is implemented in ION) has an advanced retransmission timeout estimation mechanism, which is based on the provided contact times among communicating nodes and apart from the inbound and outbound queuing delays takes also into account the communication availability for timers suspension.

One factor that strongly affects protocols performance into these challenging environments is the type and frequency of the feedback from the receiver. CFDP supports the transmission of negative acknowledgments from the receiver side in four different modes; Deferred, Immediate, Prompted and Asynchronous. LTP's main acknowledgment strategy is similar to CFDP deferred mode, but with enhanced flexibility. Whenever LTP is used as a convergence layer protocol for BP, it can aggregate continuous sequence of bundles destined to the same LTP engine and which require the same level of delivery assurance (i.e., the same LTP colour) into single blocks in order to control the acknowledgment traffic. The last segment of a red coloured block (EORP segment) is marked as a checkpoint triggering the transmission of a reception report upon receiving that segment. Additionally to this deferred-like mode, LTP allows for incremental data retransmission either by marking any red-part segment prior to the EORP as a checkpoint or by transmitting asynchronous reception reports from the receiver side. From the above discussion we conclude that unlike CFDP, LTP offers enhanced flexibility regarding retransmission techniques.

Both CFDP and LTP allow for variable segment size depending on the MTU of the underlying network and do not incorporate any mechanism for congestion and flow control.

Departing from the above discussion, we highlight the advantages of an integration of CFDP into ION and provide all the necessary parameters that allow for an efficient interoperation among CFDP, BP and LTP. Since LTP presents enhanced retransmission ability, it is evident that the reliability mechanisms of BP/LTP should be maintained, while the transport layer functions of CFDP can be disabled (i.e., by using CFDP in unacknowledged mode).

In order to run CFDP on top of the BP/LTP stack we first define how transmitted data are presented in each layer. Since BP/LTP allows for data aggregation into blocks prior to segmenting to LTP segments, choosing a relative small CFDP PDU minimizes head-of-line blocking and provides fine, application-appropriate granularity whenever multiple paths exist, while the function of data aggregation and control of the acknowledgment traffic is left to the underlying network. Thus, each CFDP PDU is encapsulated into a small-sized bundle (e.g., 1Kbyte), LTP convergence layer adapter aggregates these bundles into red coloured blocks and presents them to LTP for transmission. The number of bundles that can be aggregated into a block is limited by a maximum block size which in turn controls the maximum aggregation delay and the report traffic. Large block sizes allow for coarse-grained transmissions, but impose additional delay to partial data delivery especially when operating at high error rate environments; each block and the respective application PDUs are delivered to the upper-layers only upon the reception of the entire block. Concluding, smaller block sizes allow for faster partial data delivery to CFDP and an acknowledged mode similar to the CFDP immediate mode, but increase report traffic and impose higher complexity at both sender and receiver due to the higher number of retransmission timers that should be managed.

IV. EVALUATION

The performance of the CFDP-ION integration was evaluated using COMNET Lab's space-oriented communication testbed [9]. The testbed consists of several nodes implementing a range of space protocols, such as ESA ESTEC's CFDP implementation, both Bundle Protocol implementations (ION and DTN2), LTP, AMS, etc., and can emulate complex space communication scenarios.

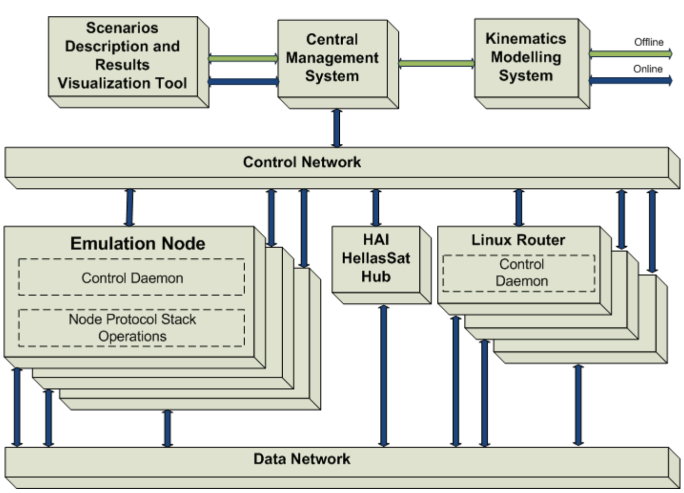

Figure 2. DTN Testbed Architecture

Both CFDP and ION are able to transfer data reliably; therefore we disabled CFDP's reliability by using it in unacknowledged mode. Reliability was assured by the underlying ION's convergence layer protocol.

A. *Terrestrial and geostationary links*

In order to evaluate the performance of CFDP over delay tolerant networks in terrestrial environments with low propagation delays and space links up to the distance of geostatic satellites we use 50 - 300 ms propagation delay over BP/TCP. At higher propagation delays, TCP performance is highly decreased, the protocol malfunctions due to its timers, as its purpose is the communication over terrestrial links with propagation delay at the order of some milliseconds. The packet loss rates vary from 0 to 10 percent, while the bandwidth is asymmetrical, using a 1Mbps downlink and a 256 Kbps uplink to deliver a file of 10 MB in a 2-node topology. The size of each CFDP PDU equals to 1024 bytes.

In Figure 3 we present the results using TCP as the underlying reliable transport protocol. As expected, file delivery time increases with the packet loss rate. Furthermore, we observe that for higher propagation delays the impact on file delivery time is greater. For packet loss rates lower than 1% the performance of CFDP in not affected by the propagation delay variance. The results show that CFDP can be used under these conditions quite successfully, especially if we consider that CFDP as a stand-alone application is designed for much higher delays, in the order of seconds.

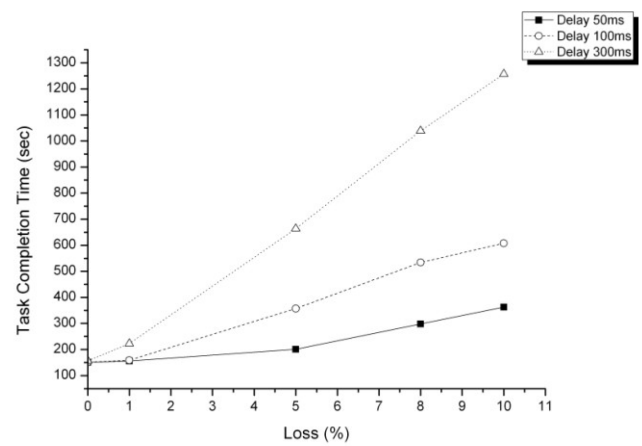

Figure 3. CFDP over ION TCP

B. *Deep space links*

Space rovers and satellites, are able to send data at rather high speed, while on the contrary earth to space links are characterized by lower speed. On that basis we emulate space links using asymmetric bandwidth and variable error rates at forward and reverse link. Specifically we use a 1 Mbps downlink and a 256 Kbps uplink, with a 30 seconds propagation delay in each and packet loss rates 0, 0.1, 1, 5 and 10 percent in both downlink and uplink. We evaluate CFDP in unreliable mode over ION with LTP protocol selected for reliability and compare it to CFDP in reliable mode over UDP. In order for LTP's and CFDP's retransmission mechanism to be comparable, CFDP is set in immediate NACK mode.

The LTP implementation we use is designed to operate in deep space links where data loss due to corruption (radiation, limited transmission power) is generally minimized by heavy forward error correction coding at the link level. Moreover, it is designed for moderate data rates and the high data rates that are used in our experiments result in high burst rates that cause the malfunction of the receivers' retransmission timers. In our scenarios, LTP could not function properly at high packet loss rates; therefore, we present LTP results only up to 1 percent packet loss.

In Figure 4 we show the results for a 5 MB file transfer in CFDP PDUs of 1024 bytes over UDP and BP/LTP in a 2-node topology. When BP/LTP are used as the convergence layer protocols of CFDP, file delivery time is better to UDP for packet loss rates up to 1 percent. Nevertheless, CFDP over UDP can operate in much higher packet error rates. The results for the 10 MB file transfers and for different packet loss rates presented in Figure 4 verify the aforementioned behavior. Each file is segmented into CFDP PDUs of 1024 bytes as before. Higher number of packets causes much higher burst rates, and hence more retransmissions. CFDP over BP/LTP performs better than UDP for packet loss rates up to 0.1 percent, but when packet loss rate increases LTP performance decreases. On the contrary, CFDP over UDP transfers of 5 and 10 MB files appear to be smoother than those of LTP for the same packet error rates as shown in Figure 4.

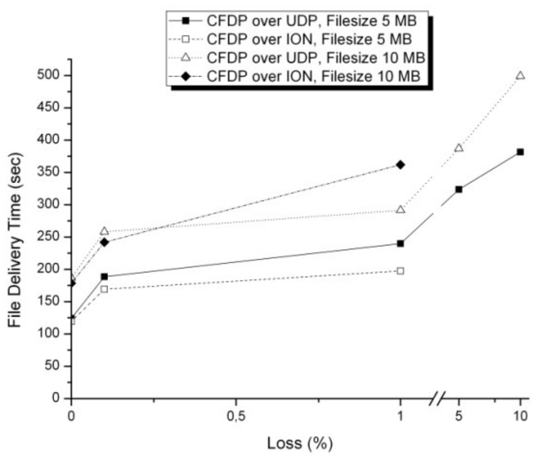

Figure 4. CFDP over ION LTP

V. CONCLUSION

In this paper we studied the integration of CFDP, the CCSDS file transfer protocol for space communications, within the DTN architecture and we evaluated the potential advantages of this integration. Each component complements the functionality of the other, rendering this combination a promising solution for file delivery. Our preliminary results show that this integration is feasible. CFDP can be used as a space file transfer application, relying on LTP's reliability and Bundle Protocol's routing scheme. However, further experiments are still needed in order to identify potential trade-offs, determine those space environments at which this integration can be beneficial and the mechanisms required to maximize performance. Future versions of LTP, designed for higher data loss rates and improved retransmission mechanism would perform much better, resulting in a successful integration of CFDP and Bundle protocols.

REFERENCES

[1] "Space Packet Protocol". Recommendation for Space Data System Standards, CCSDS 133.0-B-1. Blue Book. Issue 1. September 2003.

[2] "CCSDS File Delivery Protocol (CFDP) Part 1: Introduction and Overview", Informational Report, CCSDS 720.1-G-3. Green Book. Issue 3. April 2007.

[3] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, "Delay-Tolerant Network Architecture", RFC 4838, April 2007.

[4] K. Scott and S. Burleigh, "Bundle Protocol Specification", RFC 5050, November 2007.

[5] Jet Propulsion Laboratory, "ION: Interplanetary Overlay Network", available at https://ion.ocp.ohiou.edu .

[6] M. Ramadas, S. Burleigh, and S. Farrell, "Licklider Transmission Protocol - Specification", IETF RFC 5326, experimental, September 2008, URL: http://www.ietf.org/rfc/rfc5326.txt.

[7] Scott Burleigh, "Dynamic Routing for Delay-Tolerant Networking in Space", Flight Operations SpaceOps 2008 Conference.

[8] The Delay-Tolerant Networking Research Group (DTNRG): http://www.dtnrg.org/ .

[9] C. Samaras, I. Komnios, D. Sotirios, E. Koutsogiannis, V. Tsaoussidis, G. Papastergiou and N. Peccia, "Extending Internet into Space – ESA DTN Testbed Implementation and Evaluation", 1st International Conference on Mobile Lightweight Wireless Systems, Workshop on Research activities funded by ESA in Greece, Athens, Greece, May 18-20, 2009.