Study on CFDP and DTN Architectures for ESA Space Missions

Felix Flentge
Ground Systems Engineering Department
ESA / ESOC
Darmstadt, Germany
felix.flentge@esa.int

***Abstract*—Upcoming file-based operations concepts and more complex communication topologies for space missions require modifications and extensions to the 'traditional' space and ground communication architectures. This paper reports on a study initiated by the European Space Agency (ESA) to analyse the suitability of the CCSDS File Delivery Protocol (CFDP) and Delay Tolerant Networking (DTN) architectures for future ESA missions. Starting from an analysis of potential future missions, generic mission scenarios involving complex communication topologies have been studied and corresponding communication requirements have been defined. Based on these generic mission scenarios and the communication requirements two reference architectures have been designed: one using CFDP on top of already deployed protocols and one combined CFDP/DTN architecture. A simulation environment to evaluate the reference architectures has been created and various communication scenarios have been evaluated. The paper introduces both reference architectures and presents some results of the simulation activities. These results and further analysis of the reference architectures lead to the conclusion that both architectures are quite similar in terms of performance and can satisfy most requirements. However, in the short to medium-term time frame CFDP without DTN seems to provide the easier way to adoption since it is conceptually simpler and more mature while the additional features of DTN may only be required in the long-term.**

Keywords - Bundle Protocol (BP); CCSDS File Delivery Protocol (CFDP); Delay Tolerant Networking (DTN); Licklider Transmission Protocol (LTP)

I. INTRODUTION

The trends towards more file-based operation concepts for space missions and increased complexity of space and ground communication topologies (e.g., data relays in space) has lead ESOC (the European Space Operations Centre) to initiate a study analysing how CFDP and DTN could be utilised in future ESA missions. Currently, support for file transfers and data relaying is mainly implemented by 'private' means for each mission. In conjunction with ongoing work on file-based operations at ESOC [1], generic communication requirements to support file-based operations and complex communication topologies have been defined.

The Consultative Committee for Space Data Systems (CCSDS) has standardised a CCSDS File Delivery Protocol (CFDP) [2] and is in the process of standardising protocols for Delay Tolerant Networking (DTN) in space [3]. CFDP is a file delivery protocol allowing file transfers over multiple hops and taking the specific space environment (like line-of-sight disruptions, long delays, etc) into account. DTN is an architecture for internetworking of networks that may be separated by disruption or delays. For space, it is typically implemented by using the Bundle Protocol (BP) [4] and underlying convergence layer protocols, like the Licklider Transmission Protocol (LTP) [5] or Proximity-1 [6]. Based on the generic communication requirements and the mission scenarios two reference communication architectures involving ground and space segment have been designed: one using CFDP on top of already deployed protocols and one combined CFDP/DTN architecture. In the former architecture store-and-forward features and retransmission capabilities of CFDP are utilised to provide multi-hop file transfer and reliability. In the latter architecture these functionalities are provided by BP and LTP, and CFDP just provides file transfer. Both reference architectures have been analysed and simulated using the ESOC Ground System Test and Validation Infrastructure (GSTVi) [7].

The paper starts with introducing the mission scenarios, communication requirements and the relevant standards. The reference architectures are presented and results from the simulation activities are reported. The paper finishes with some conclusions concerning the use of CFDP and DTN for future ESA missions and points to areas for future work.

II. MISSION SCENARIOS & COMMUNICATION REQUIREMENTS

In order to define generic communication requirements ongoing and planned ESA missions have been analysed in terms of communication architectures and topologies. Astronomy, Earth observation and planetary missions involving planetary landers have been taken into account. A 'Complex Mission Topology' has been derived from this analysis and includes all elements of the relevant missions:

- A complex space segment containing spacecrafts potentially owned by different space agencies that may be used as data relays as well as a landed asset on a planetary surface.

 ISBN: 978-1-61208-128-1

- An interoperating ground segment with ground stations owned by different agencies, (interoperating) Mission Control Centres (MCC), Lander Control Centres (LCC) and User Support and Operations Control Centres (USOC).
- Bi-directional asymmetric and uni-directional potentially disrupted and delayed communication links between the various elements.

The Complex Mission Scenario has formed the basis for the definition of the CFDP and the CFDP/DTN reference architectures. Detailed communication requirements have been defined based on the analysed missions and in conjunction with a working group on file-based operations at ESOC [1]. These requirements are related to different topics:

- **Communication Environment**: The reference architecture has to take the special characteristics of space communication into account. This includes long communication delays, low signal-to-noise ratios, high bandwidth/delay products as well as predictable (orbits, planetary rotation) and unpredictable disruptions (e.g. solar environment).
- **Communication Services**: Two types of co-existing communication services should be supported:
    - **File Services**: The transport of large data structures that are self-contained and persisted at source and target destination in file systems.
    - **Message Services**: Usually smaller data structures that are atomic and immediate in nature from operations point of view.
- **Quality of Service** aspects including completeness, error-free and in-sequence delivery may have to be regarded. Priority and pre-emption of certain data should be ensured, notification of end-to-end delivery and session control have to be provided.
- **Data Management** (of data in transfer): Information related to data transfer has to be available to the operators and data management operations (deletion of data in transfer, queue re-ordering, changing priorities, etc) should be possible to give the operators full control of all data transfer.
- **Security and Safety** including authentication, integrity and confidentiality services has to be taken into account. Of particular interest are also mission safety concerns. The Mission Control Centre (MCC) should have full visibility and control over the data uploaded to a spacecraft.
- **Routing** is expected to be static and planned based on link availability. Recovery from failure is likely to take the form of using pre-configured backup routes rather than dynamic route discovery.
- **Interoperability** at protocol level and cross-support between different agencies must be possible at various points, e.g., between MCC and ground station networks, ground stations and data relays in space or data relays and landed assets.

III. STANDARDS AND PROTOCOLS

The reference architectures take available and upcoming communication standards into account to facilitate future deployments and cross-support between different agencies. The following standards are used:

- **CFDP**: CFDP is used for bi-directional file transfer between various elements in the architecture [2]. CFDP includes different classes for reliable (Class 1) and unreliable (Class 2) file transfer as well as special features for relaying file transfers through intermediate waypoints (Store-and-Forward Overlay; CFDP Class 3 for unreliable and CFDP Class 4 for reliable multi-hop transfer). CFDP also contains a set of file manipulation primitives for managing remote file stores by operations such as file or directory creation, deletion and copying.
- **DTN**: DTN (Delay or Disruption Tolerant Networking) is an architecture for internetworking between separated (e.g., by delays or disruption) networks [3]. This is implemented by the use of the Bundle Protocol (BP) for the necessary store-and-forward capabilities and Convergence Layer Protocols for providing data transport [4]. For terrestrial networks this may be provided by UDP or TCP while for space links LTP (Licklider Transmission Protocol) may be chosen [5]. LTP is based on a negative acknowledgment scheme similar to the one used by CFDP and allows marking parts of the data for reliable (red mode) and for unreliable (green mode) data transport.
- **CCSDS Data Link Layer Protocols**: For the space links the usual CCSDS Data Link Layer Protocols are used. These include Packet Telemetry (TM) and Telecommand (TC), Advanced Orbiting Systems (AOS) Space Data Link and Proximity-1 for the link between a landed asset and an orbiter [6].
- **Encapsulation Packet**: The CCSDS Encapsulation Packet provides the means to multiplex packets from different user protocols (like space packets or CFDP) into the space link [6].
- **SLE**: CCSDS Space Link Extensions services are used to extend the services offered by the Data Link Layer Protocols from the ground station to the service user at the mission control centre [6].

IV. COMMUNICATION REFERENCE ARCHITECTURES

Based on the communication requirements and the available standards and protocols two reference architectures, one applying CFDP on top of ‘traditional’ protocols and one applying CFDP on top of DTN protocols, have been designed. Both reference architectures include the same key elements:

- A **Lander** on a planetary surface.
- An **Orbiter** orbiting this planet and used as a data relay for the lander.
- A **Ground Station (Network) (GS)** on Earth for communication with the orbiter and the lander.

 ISBN: 978-1-61208-128-1

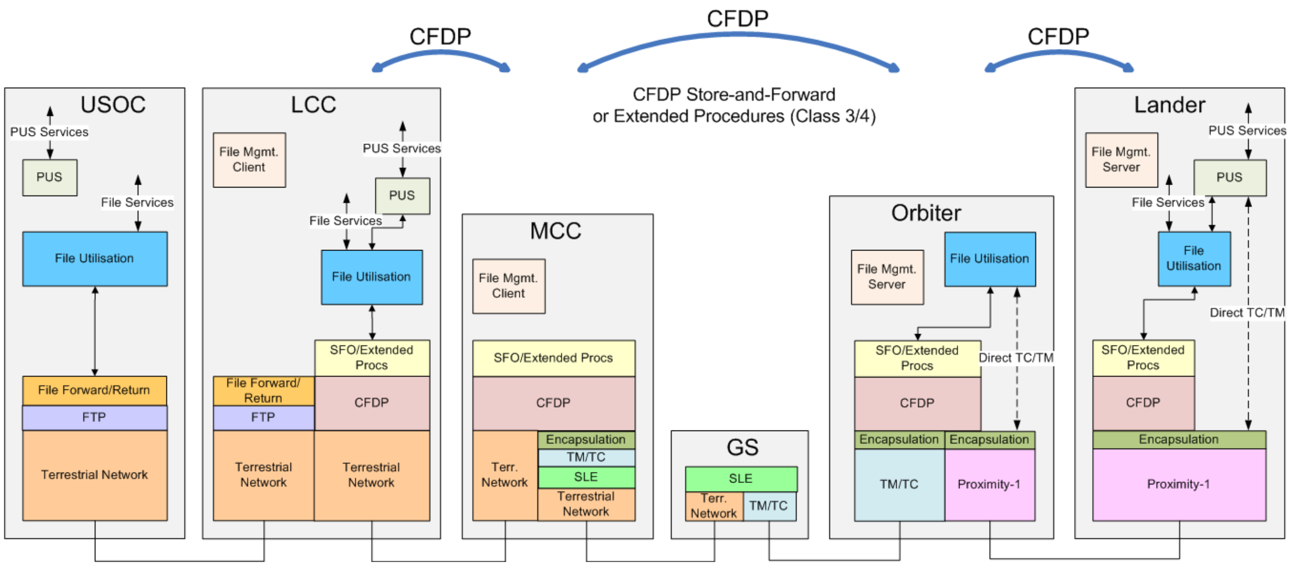

Figure 1. CFDP Reference Architecture.

- A **Mission Control Centre (MCC)** for monitoring and control of the orbiter.
- A **Lander Control Centre (LCC)** to monitor and control the lander. Data transfer using the orbiter is through the MCC.
- A **User Support and Operations Centre (USOC)** to monitor and control payloads onboard the lander. This is done by utilising the ESA Packet Utilisation Standard (PUS) [8] and a (not-yet existing) File Utilisation Standard (FUS). Data transfer is always through the LCC.

A. CFDP Reference Architecture

The CFDP Reference Architecture is shown in Figure 1. The main features of the reference architecture are:

- Use of a **conventional file transfer protocol** (like FTP) between USOC and LCC potentially enhanced by a File Forward/Return Service. A conventional file transfer protocol is chosen instead of CFDP as certain functions would have to be performed by the LCC taking mission-wide factors into account :
    - A **mission safety firewall** function at the file or packet utilisation level, which can check semantics of the file uplink/forward data against, for instance, permissions, resource allocation, etc.
    - A **pro-active fragmentation** function that, in the event of the use of multiple orbiter relays, breaks files into fragments that allow management of the transfer of each fragment in earth-orbiter and orbiter-lander contact periods.
    - A **remote file management** client, which exists to satisfy requirements for remote file fragmentation, queue reordering, status reporting and file transfer pre-emption.
- Use of **CFDP Class 1/Class 2 and Store-and-Forward Overlay (or CFDP Class 3/Class 4)** between LCC, MCC, orbiter and lander.
- Use of **SLE between MCC and Ground Station (Network)**. As the MCC – Ground Station Links can be assumed to be continuously available there seems to be no need to place CFDP inside the Ground Station as fragmentation and scheduling could be performed more easily from the MCC, in particular in the case multiple ground stations are used and files would have to be distributed to different ground stations. However, especially with the downlink of scientific data there are cases where it makes sense to terminate the CFDP traffic in the ground station. These cases have not been studied in detail, yet.
- Use of **CFDP or direct TC between orbiter and lander on top of Proximity-1**. The orbiter should provide the option to function as a CFDP intermediate waypoint (store-and-forward overlay or extended procedures) and the option to extract TC from a file and send them directly to the lander (e.g., for emergency commanding).

Please note that message services can be provided in parallel to the file transfer using the PUS standard.

B. DTN Reference Architecture

For the CFDP/DTN architecture 'pure' file transfer functionalities are still provided by CFDP while store-and-forward and reliability are provided by BP and LTP or Proximity-1 as underlying protocols. The CFDP/DTN Reference Architecture is shown in Figure 2. The main features are:

- Use of **CFDP Class 1 (unreliable transfer) between LCC and orbiter (direct TC) or lander**. Reliability will be provided by the underlying protocols and store-and-forward is realised by BP.

 ISBN: 978-1-61208-128-1

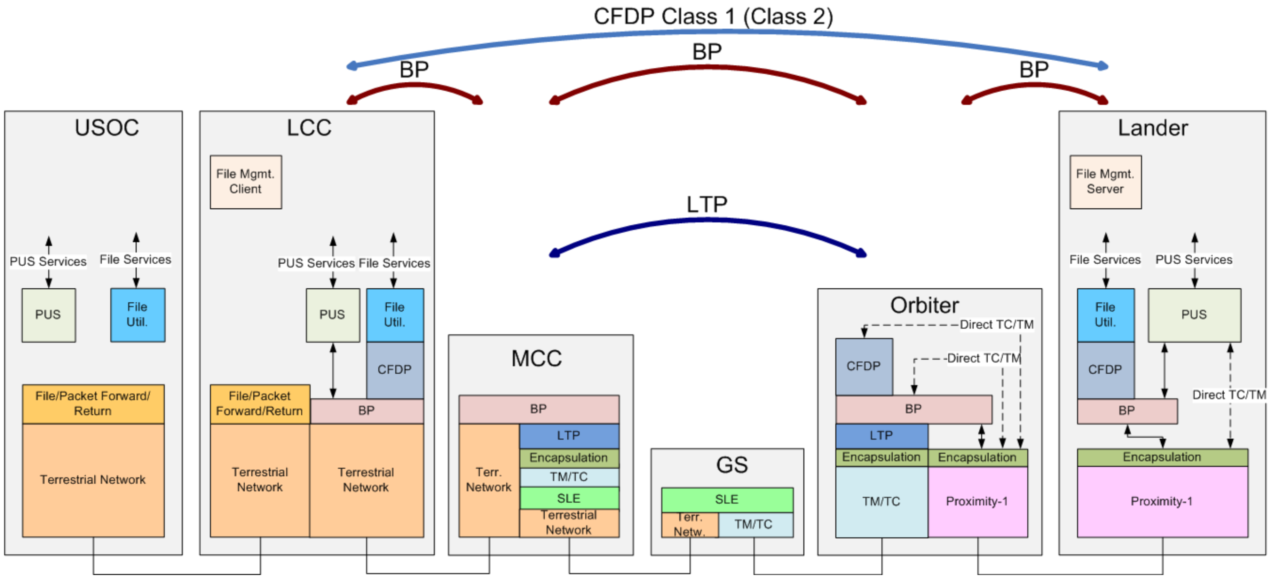

Figure 2. CFDP/DTN Reference Architecture.

As BP is a transport layer protocol additional means for end-to-end acknowledgement for a complete file transfer may be needed. This could be provided by using bundle status reports, using CFDP class 2 (without really needing retransmission) or a yet-to-be-defined CFDP Class 1a without retransmissions but with end-to-end confirmation.

- Use of **SLE between MCC and Ground Station (Network)**. As the MCC – Ground Station Links can be assumed to be continuously available there seems to be no need to apply the store-and-forward of BP in the ground station. However, under some circumstances (downlink with routing of bundles to different users, low link capacity between GS and MCC) supporting BP and providing bundle storage in the ground station is desirable.
- Use of **BP over LTP and CCSDS Encapsulation Packets between MCC and Orbiter** with LTP providing (selective) reliability and the encapsulation packet providing multiplexing of different protocols if necessary.
- Use of **BP or direct TC on top of proximity-1 between orbiter and lander**. Again, a direct TC capability should be provided by the orbiter by either extracting TC packets from bundles or from files.

In this case message services can be provided by PUS or the upcoming Asynchronous Message Service [6].

V. ARCHITECTURE SIMULATION

In order to validate the proposed reference architectures a simulator based on ESOC's GSTVi [7] and the SIMSAT simulation framework has been created. For initial tests SCOS 2000 has been connected as Mission Control System to initiate file transactions, send TC and receive TM packets. For more systematic and automated tests a traffic load generator component has been used. ESOC's CFDP Entity implementation for the ground segment [9], the DTN2 implementation of the BP from the Delay Tolerant Network Research Group and the LTPlib from Trinity College Dublin for LTP have been used to simulate the necessary protocols. As no Proximity-1 implementation has been available, only a basic Proximity-1 emulation based on UDP has been created.

A. A. Mission Test Scenarios

Many simulations with different mission configurations have been performed. In this paper we concentrate on:

- '**Relay Mission with CFDP Only'** including the MCC, a single ground station, an orbiter and a lander.
- '**Relay Mission with CFDP over DTN**' including the same elements as the 'Relay Mission with CFDP Only' mission configuration but using CFDP on top of BP and LTP.

Acknowledged and unacknowledged file transfers putting files to the simulated lander and getting files from the lander (using CFDP proxy put requests) have been evaluated. Getting a file involves both - sending a request to the lander and the actual transmission of the file to Earth. Two 'error conditions' have been tested:

- **Best Case**: Continuous end-to-end link without QoS errors, i.e., no retransmissions occur.
- **Dropped packets**: The 4th TM packet and the 4th TC packet that are send from/to ground are dropped. In acknowledged modes, this will lead to retransmissions.

Typical bandwidths have been selected: 10 kbps for the earth-orbiter link and 500 kbps for the orbiter to lander link. A large latency between Earth and orbiter (1200 sec) and a small latency between orbiter and lander (1 sec) have been used. CFDP PDU sizes have been chosen to be completely carried in a PDU of the most restrictive underlying protocol (220 byte / 1024 byte for 'CFDP only' uplink / downlink on earth-orbiter link; 10240 byte for 'CFDP/DTN' for uplink and downlink on earth-orbiter link). The experiments have been performed with different file sizes (1k, 10k, 100k).

B. Simulation Results

During the simulations it has been detected that getting files from the lander using a CFDP proxy put request could not be executed for the relay missions with CFDP Store-and-Forward Overlay (SFO). The reason is that SFO currently cannot carry a proxy request to the destination waypoint. This has been reported to the CCSDS Working Group.

Furthermore, 100k file transfers with CFDP/DTN have not been possible with the UDP-emulated Proximity-1 link since files have been put in a single bundle and bundles exceeding the maximum size of UDP datagrams cannot be send over the UDP convergence layer.

1) Unacknowledged File Transfers

For unacknowledged file transfers CFDP Class 3 was compared with CFDP Class 1 SFO and CFDP Class 1/BP/LTP green mode.

a) Putting files to the spacecraft – Best Case

The time needed for sending a file from ground to the lander is only slightly above the One Way Light Time (OWLT) for all protocol configurations with very little differences. As shown in Figure 3, the protocol overhead on the orbiter to lander link is quite high for small files (around 15%) but drops significantly for larger file sizes. For large file sizes there is a considerable overhead for CFDP class 3 compared to the other protocol configurations. This can be attributed to the fact that CFDP Class 3 because of the CFDP PDU forwarding mechanism uses a smaller CFDP PDU size on this link (220 bytes) compared to the other protocol configurations. For example, in the 10k file size case 49 packets are send, compared to respectively 8 and 6 packets for Class 1 SFO and Class 1/BP/LTP green.

b) Getting files from the spacecraft – Best Case

As explained above, the CFDP SFO does currently not allow initiating a proxy get operation over hops, so just CFDP Class 3 and CFDP Class 1/BP/LTP green have been compared. Transaction durations are for both configuration about 2 x OWLT (1 OWLT to initiate the put request to downlink the file + 1 OWLT to downlink the file).For the 1k file there is about 15% protocol overhead for CFDP Class 3 and about 23% protocol overhead for CFDP Class 1/BP/LTP green on the lander to orbiter link. For 10k files, the protocol overhead is only about 2.5% for both configurations.

2) Acknowledged File Transfers

For acknowledged file transfers CFDP Class 2 SFO and CFDP Class 4 were compared with CFDP Class 1/BP/LTP red and CFDP Class 2/BP/LTP green (to see whether reliability should be provided by CFDP or LTP). Dropping of packets has only been investigated with 10k files.

a) Putting files to the spacecraft – Best Case

Transaction durations are about 3 OWLT for all protocol configurations but Class 2 SFO, which needs about 4 OWLT. Three OWLT are needed for sending the file to the lander, sending a notification that the file transfer has finished back to ground and getting an acknowledgment for this notification. Class 2 SFO needs more time because the file transaction between ground and the orbiter has to be completed (3 OWLT) before the file is sent to the lander and a SFO Report is sent back from the lander to the ground (1 OWLT). As usual, protocol overheads are quite high for 1k files (15% to 24 % with CFDP Class 2/BP/LTP green having the highest overhead) but are small for larger files (5% for CFDP Class 4 and around 2.5% for the rest).

b) Putting files to the spacecraft – Dropped Packet

The whole picture changes as soon as TC/TM packets are dropped (see Figure 4). CFDP Class 4 has the shortest duration (4 OWLT), with CFDP Class 1/BP/LTP red and CFDP Class 2 SFO taking 50% longer. CFDP Class 2/BP/LTP green takes twice as long as CFDP Class 4. The good performance for CFDP Class 4 is due to the fact that the retransmission request for the lost TM is send back to Earth immediately and is not affected by the dropped TC packet (as only the 4th packet is dropped). For CFDP Class 2 SFO and CFDP Class 1/BP/LTP red, the completion of the file transfer is affected by the dropped TM and the dropped TC and taking 2 OWLT more. CFDP Class 2/BP/LTP green detects the missing TM packet only at the lander and is affected by the dropped TC. This leads to a retransmission of the whole bundle and a very long duration (8 OWLT).

The protocol overheads for the 10k file are similar to the results without dropped packages with the exception of CFDP Class 2/BP/LTP green that has an overhead of 100% that reflects the resending of the whole bundle.

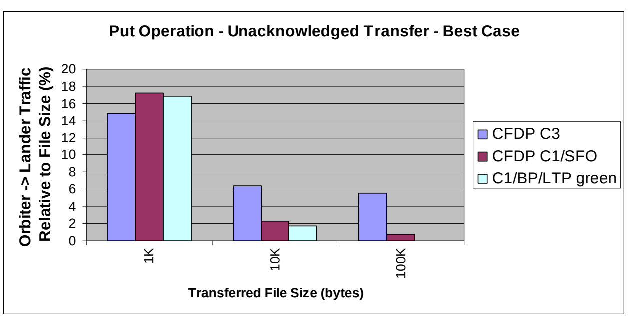

Figure 3. Comparison of protocol overhead on the orbiter to lander link by different communication architectures to an unacknowledged file transfer to a lander in best case error conditions.

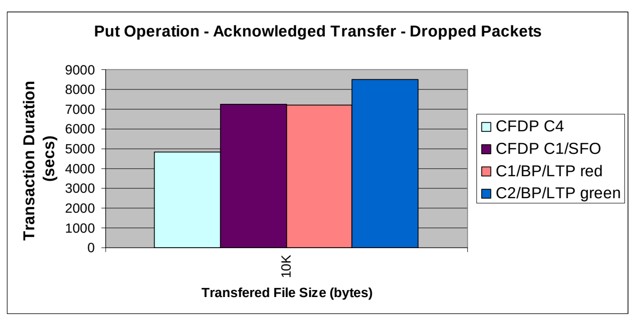

Figure 4. Comparison of transaction durations by different communication architectures to an acknowledged file transfer to a lander in dropped packets error conditions.

c) Getting files from the spacecraft – Best Case

The transaction durations for getting a file from the spacecraft (i.e., proxy put request + acknowledged downlink of the file) takes for all protocol configurations about 4 OWLT (1 OWLT to initiate the downlink + 3 OWLT for the downlink; SFO is not possible in this case). Protocol overhead is high to very high for 1k files (23% CFDP Class 1/BP/LTP red, 32% CFDP Class 2/BP/LTP green, 50% CFDP Class 4) but comparable to the overhead in other scenarios for bigger file sizes (2.3% to 3.5%) with a larger overhead for CFDP Class 4 (5.4%), which can again be attributed to the smaller PDU size as explained for CFDP Class 3 in the unacknowledged case.

d) Getting files from the spacecraft – Dropped Packet

For acknowledged transfer with dropped packets CFDP Class 4 and CFDP Class 1/ BP/LTP red need about 8 OWLT. This is because a TC and a TM packet are dropped, causes retransmissions due to negative acknowledgements and inactivity timeouts. In this scenario, CFDP Class 2/BP/LTP green could not deliver the complete file because LTP is used in unreliable mode and one TC packet containing an LTP segment is dropped. At the receiving end the bundle is reassembled but part of the data is missing, leads to an incomplete file when the bundle is passed to CFDP. As CRC has not been used for CFDP PDUs in this case, the PDU is accepted but a file checksum error is raised. Looking at protocol overhead with 10k files reveals again a larger overhead for CFDP Class 4 (8%) compared to CFDP Class 1/BP/LTP green (2%) caused by the smaller CFDP PDU on the lander to orbiter link for CFDP Class 4.

VI. Conclusion and Future Work

Both, the CFDP only and the CFDP/DTN reference architectures, provide solutions to the provision of message and file services over noisy, long-delay, disrupted channels as found in typical ESA missions. The preliminary conclusion is that, once the timeouts were tuned to those appropriate for the mission scenario, there is little to choose between any of the communication architectures in terms of transaction duration and protocol overheads. For retransmission of lost packages hop-by-hop retransmission (CFDP Class 4, SFO, LTP Red Mode) is generally preferable over end-to-end retransmissions (CFDP Class 2 over BP). SFO has a slight disadvantage in terms of transaction duration as all of the file must be transferred to and stored on the relay before it can be forwarded. Furthermore, CFDP SFO has currently some conceptual problems with proxy put requests over multiple waypoints.

However, for specific missions, further simulation with more realistic link characteristics and an optimisation of the protocol configurations as well as formal analysis is needed to compare potential architectures more realistically and to understand potential trade-offs. For example, on the one hand CFDP lacks security primitives, and arguably includes too many “layers” in a single specification while DTN is architecturally cleaner and is not limited to file transfers. On the other hand, some aspects of DTN are still subject to active research (network management, routing, security aspects for BP) and there are a number of operational uncertainties that need to be studied before one would mandate its use in future ESA missions. In particular, the network management aspects of DTN are not yet well-developed and may require significant additions to the BP before DTN can meet the current operational requirements.

Apart from these operational aspects, for the missions and scenarios analysed in the study, the bulk of the requirements can be satisfied by adoption of existing CCSDS recommendations including CFDP or by a CFDP/DTN architecture. However, some of the additional features that DTN supports are not yet required for upcoming and planned ESA missions. From the operational point of view there may be difficulties in accepting the new model of control that is implied with a DTN based architecture and further research and development into DTN may be necessary to establish that DTN can meet these operational requirements. In addition, CFDP has been already standardised by CCSDS and ESA implementations for the ground and space exist while DTN-related protocols are still in the CCSDS standardisation process. So, in the short to medium-term time frame CFDP without DTN seems to provide the easier way to adoption since it is conceptually simpler and more mature while the additional features of DTN like dynamic routing or reactive fragmentation may only be required in the long-term. DTN does appear to provide a solution in the long-term when more complex networks of interoperating assets in space and on ground (Earth and other planets) exist.

Acknowledgment

The CFDP/DTN Study has been performed by SciSys, UK in cooperation with Trinity College Dublin, Ireland and Keltik Ltd., UK. Without their excellent work and expertise it would not have been possible to achieve the high quality results reported in this paper.

References

[1] M. Pecchioli, E. Montagnon, G.P. Calzolari, F. Flentge, and A. Kowalczyk, “Towards File Based Operations”, AIAA 2010-2190, Electronic Proceedings for SpaceOps 2010.

[2] CCSDS 727.0-B-4. CCSDS File Delivery Protocol (CFDP), Recommended Standard, Blue Book, January 2007.

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

[4] K. Scott and S. Burleigh, “Bundle Protocol Specification”, Network Working Group, RFC 5050, November 2007.

[5] M. Ramadas, S. Burleigh, and S. Farrell, “Licklider Transmission Protocol – Specification”, Network Working Group, RFC 5326, September 2008.

[6] Consultative Committee for Space Data Systems (CCSDS), http://public.ccsds.org, access date: 31 January 2011.

[7] M. Goetzelmann, N. Lindmann, P. Ellsiepen, C. Laroque, M. Niezette, and A. Walsh, “Ground Systems Test and Validation Infrastructure – from Concept to Implementation”, AIAA 2006-5540, Electronic Proceedings for SpaceOps 2006.

[8] ECSS-E-70-41A. Ground Systems and Operations – Telemetry and Telecommand Packet Utilization, January 2003.

[9] F. Flentge and M. Jordan, “Design and Implementation of CCSDS File Delivery Protocol for ESA Ground Segment”, ESA SP-669 (CD-ROM), Proceedings of DASIA 2009.

 ISBN: 978-1-61208-128-1