Options for Repeated Transmissions of TC Transfer Frames

1. Introduction

This document considers the options for the specification of systematic retransmissions of TC Transfer Frames, under the control of a new parameter in the CCSDS telecommand protocol.

1.1 Related CCSDS standards

The following CCSDS recommended standards (Blue Books) are relevant to the telecommand retransmissions discussed in this document:

- TC Synchronization and Channel Coding, CCSDS 231.0-B-1, September 2003
- TC Space Data Link Protocol, CCSDS 232.0-B-1, September 2003
- Communications Operation Procedure-1, CCSDS 232.1-B-1, September 2003

1.2 COP-1

The Communication Operations Procedure 1 (COP-1) is in the Transfer Sublayer of the Data Link Layer of the CCSDS telecommand communication protocol stack. It provides two data transfer services called the sequence-controlled service and the expedited service. COP-1 uses three types of Telecommand Transfer Frames, which are indicated by flags in the frame header:

- Type-AD frames carry data for the sequence-controlled service. During operation of the service, there can be many type-AD frames outstanding (unacknowledged). This document is principally concerned with the systematic retransmission of type-AD frames.
- Type-BC frames carry control commands for the sequence-controlled service. They can be used when starting a session for the sequence-controlled service, as part of the process for ensuring that the sender and receiver are correctly synchronized for their use of sequence numbers. There is never more than one type-BC frame outstanding. The application of the systematic retransmission to type-BC frames is also considered.
- Type-BD frames carry data for the expedited service. They are not retransmitted, and any new systematic retransmission mechanism that is introduced for type-AD frames should ensure that it does not result in the retransmission of type-BD frames.

1.3 Purpose of the systematic retransmissions

The systematic retransmission is intended to improve the efficiency of COP-1 for deep space missions on links with long light delays. The COP-1 protocol includes an automatic retransmission mechanism which works very well for near-Earth missions: the total round-trip delay is short so COP-1 can respond efficiently to the need for retransmission of a lost frame. For deep space missions the round-trip delay is long and the COP-1 response is not efficient. As discussed below, the existing COP-1 parameters can be manipulated to force early automatic retransmissions for a deep space mission.

For a sequence of frames ABCDEF…, the COP-1 automatic retransmission results in a cyclic transmission pattern of the form:

ABCDEF…ABCDEF…ABCDEF…

This document considers the options for adding a new function under the control of a new parameter: With the proposed new parameter, the user of the protocol would be able to force the transmission of type-AD frames a specified number of times, instead of only once. The specified number of systematic retransmissions would be used at every transmission attempt. For a sequence of frames ABCDEF…, the systematic retransmission would result in a hammer transmission pattern of the form:

AAABBBCCCDDDEEEFFF…

When frame loss errors are sporadic, the hammer pattern is an advantage: because the errors are sporadic, they are less likely to hit all the repeats of a frame. With cyclic transmission pattern and sporadic errors, successful reception of a sequence of frames depends on where the errors fall within a cycle. For example, a frame loss of an early frame in the first cycle causes the remaining frames in that cycle to be rejected. The advantage of the hammer pattern over the cycle pattern increases as the number of frames in the sequence increases.

Example of sporadic errors and cyclic transmission pattern:

Frames	1 2 3 4 5 6 7 . . . . . n-2 n-1 n	*1st cycle: Frame 3 lost. Following frames rejected*
Errors	X	
Frames	1 2 3 4 5 6 7 . . . . . n-2 n-1 n	*2nd cycle: Middle frame lost, following rejected*
Errors	X	
Frames	1 2 3 4 5 6 7 . . . . . n-2 n-1 n	*3rd cycle: Frame n-1 lost, so last two not delivered*
Errors	X	

When frame loss errors come in bursts, then the hammer pattern is vulnerable to the loss of all the repeats of a frame, so in this case the cyclic pattern could give better results. The hammer pattern is not a universal solution but it provides another tool for solving the problems of getting a complete sequence of frames on board.

1.3.1 Scenarios with sporadic frame loss

The typical scenario for the application of the proposed systematic retransmissions is a deep space mission with long light delay and with link conditions where errors are sporadic, so that sometimes one frame, or perhaps two frames, are lost from a sequence of frames.

When the hammer pattern of transmission is used in this scenario, it can increase the probability that the full sequence is received during a limited transmission session.

1.3.2 Scenarios with high probability of frame loss

There are deep-space scenarios where the probability of frame loss is exceptionally high:

- Spacecraft on the far side of Earth's orbit. In this scenario, the RF link to the spacecraft can suffer severe solar interference, leading to extreme frame loss rates. The telemetry link is also likely to be affected by the solar interference.
- Spacecraft unable to maintain correct attitude. In this scenario, the low-gain antenna at the receiving end of the telecommand link is intermittently pointing in the wrong direction. If the spacecraft is spinning, there is a periodic element to the high frame loss rate. When the spacecraft enters safe mode, the telecommand data rate is very low and telemetry transmissions have very low power and a low data rate.

For handling an emergency with a near-Earth spacecraft, the COP-1 expedited service can be sufficient, with telecommand frames managed one by one. For these deep-space scenarios, the proposed systematic retransmissions provide another means to attempt to deliver a sequence of telecommand frames, despite the very difficult conditions.

1.4 Scope of the systematic retransmissions

The systematic retransmission is in addition to the existing automatic retransmissions of the sequence-controlled service.

Any changes to support the systematic retransmission mechanism will be limited to the sending end of the telecommand communications system. The desired effect of the retransmissions relies on the existing behaviour of the receiving end. The discussion in this document is therefore only concerned with the sending end of the frames for the sequence-controlled service.

The expedited service of COP-1 has no retransmission mechanism: FOP-1 transmits each frame exactly once. The proposed systematic retransmission does not affect the expedited service.

It is very important to limit the scope of the modifications and the impact on existing standards. The specifications associated with the new parameter should be designed to be compatible with existing implementations, so that existing compliant implementations remain compliant with the modified standards. The behaviour of a modified system with the default value of the new parameter should be identical to an unmodified system.

2. Background

2.1 COP-1 automatic retransmissions

The sequence-controlled service includes a mechanism for the automatic retransmission of a sequence of TC Transfer Frames whenever one of them is not correctly received on board. The retransmission mechanism is “go-back-n”.

COP-1 uses protocol status information which is passed from FARM-1 (the receiving end of COP-1) to FOP-1 (the sending end of COP-1). The information acknowledges the receipt of frames and indicates problems, such as a gap in the received sequence. The information is contained in a communications link control word (CLCW), transmitted in a telemetry transfer frame.

FOP-1 starts the automatic retransmission of a sequence of frames when:

- FOP-1 receives a CLCW indicating that FARM-1 has detected a gap in the sequence, or
- FOP-1 does not receive acknowledgement of receipt of a frame within a given time.

In the automatic retransmission, FOP-1 applies the “go-back-n”: it retransmits the first missing frame (or the earliest unacknowledged frame) and all subsequent frames.

2.2 Interface for sequence-controlled service

The interface to FOP-1 for the sequence-controlled service can be considered as two parts:

- The interface for management of the sequence-controlled service.
- The interface for sending data units on the sequence-controlled service.

This view implies the existence of two user entities:

- The management entity.
- The user that is sending the telecommand data. In a typical application, the immediate source of the data is the Segmentation Sublayer, which in turn is sending data on behalf of layers above.

The management entity uses the management interface to start and stop the sequence-controlled service and to set its parameters, such as the length of the timeout period (called T1_Initial) and the limit on the number of automatic retransmissions (called Transmission_Limit). It also receives failure messages from FOP-1 and manages the consequences. The management of the normal case can easily be automated but safe recovery from some failures is more complex.

2.3 Forcing FOP-1 retransmissions with the currently-defined COP-1

When COP-1 was developed more than twenty years ago, the needs of deep-space missions were considered. The timeout period and transmission limit can be set to cause early automatic retransmissions and the Suspend / Resume mechanism can maintain the FOP-1 status, such as its internal queues, during a long waiting period.

In normal operation for a near-Earth mission, the value for T1_Initial is set to slightly more than the expected round trip time from a frame leaving FOP-1 to an acknowledging CLCW arriving back at FOP-1. To cause multiple transmissions of a sequence of frames for a deep space mission, T1_Initial and other parameters can be used differently. For example:

- The management entity sets T1_Initial to a value that is much shorter than the round-trip time and it sets the option to use the Suspend / Resume mechanism.

- FOP-1 transmits the frames in the sequence once as normal. FOP-1 keeps the frames in its internal Sent_Queue.

- Before any CLCWs can arrive to acknowledge receipt of the frames, the timer expires, so FOP-1 retransmits all the frames in the Sent_Queue. This repeats until the transmission limit is reached.

- In normal operation for a near-Earth mission, FOP-1 reports a failure when it reaches the transmission limit and it stops the sequence-controlled service and clears the Sent_Queue. In this example, FOP-1 suspends and it does not clear the Sent_Queue.

- At a later time, the management entity can adjust parameters and then reactivate FOP-1 with a Resume directive.

However, this requires actions by the management entity, in addition to the usual actions to start and stop the sequence-controlled service. The management entity has to select appropriate values for the FOP-1 parameters to force the early retransmissions. It also has to handle a FOP-1 suspend and make decisions about the timing for Resume.

Therefore, if a deep space mission is going to use this method of forcing the early automatic retransmissions, then the automation of the management entity is complex, or else the management relies on human intervention. The intention of the proposed new systematic retransmissions is to provide the user with the retransmissions but without the complex management.

3. Position of the systematic retransmission

FOP-1 is a procedure in the Transfer Sublayer of the Data Link Layer. The Transfer Sublayer receives a data unit from the next layer or sublayer above, typically the Segmentation Sublayer. Within the Transfer Sublayer, the data unit is called a Frame Data Unit (FDU) and it is placed in the Transfer Frame Data Field of a TC Transfer Frame.

What is the best position within the layers for the actions that implement the proposed systematic retransmissions? The options for the position fall into three obvious groups:

Option 1 - Before FOP-1

Option 2 - Inside FOP-1

Option 3 - After FOP-1

3.1 Option 1: Before FOP-1

Could the systematic retransmission be implemented before FOP-1?

Section 2.4 above discusses how to cause multiple transmissions by setting different values for FOP-1 variables, such as the timer. This is complex to automate and is not the best way to implement the systematic retransmissions of frames.

The proposed systematic retransmission of frames depends on existing behaviour in FARM-1 at the receiving end to reject duplicate frames, and this depends on the frame sequence numbers. The frame sequence numbers are set by FOP-1, so this excludes any mechanism for submitting the same FDU multiple times to FOP-1.

Note: It would not be completely impossible to manipulate the FOP-1 sequence number from outside. For example, a method to repeat the submission of a sequence of FDUs could include terminating FOP-1, setting its internal sequence number and then initialising it again. However, this would be extremely complex and would need more management than the methods described in section 2.4 above. To produce a hammer pattern of retransmissions would need even more intervention.

3.2 Option 2: Inside FOP-1

Clearly, the systematic retransmission can be implemented inside FOP-1. How can it be done effectively and efficiently, but with a minimum of modifications to the existing FOP-1 specification? Two different ways to implement the systematic retransmission are considered here:

- by making changes to the handling of the Sent_Queue (see section 3.2.1)
- by using the AD_Accept and BC_Accept events from the Lower Procedures (see section 3.2.2)

If either of these two is selected, then there will also be a need for a new FOP-1 parameter (see section 3.2.3).

3.2.1 Systematic retransmission using the Sent_Queue

A systematic retransmission with a “cycle” pattern:

ABCD ABCD ABCD ABCD …

could be implemented for type-AD frames using the Sent_Queue. As there is never more than one type-BC frame on the Sent_Queue, the retransmission pattern for type-BC frames can only be

AAAA…

Currently, the Sent_Queue has a To_Be_Retransmitted_Flag for each frame on the queue, which is used for the FOP-1 automatic retransmissions. This flag could be replaced or expanded by a new To_Be_Retransmitted_Count for each frame on the queue, to control the systematic retransmissions.

This option appears to be more complicated than the one described in section 3.2.2. Also, it does not add the capability for the hammer pattern of retransmissions discussed in section 1.3.

3.2.2 Systematic retransmission using AD_Accept and BC_Accept events

A systematic retransmission with a “hammer” pattern:

AAAA BBBB CCCC DDDD …

could be implemented for type-AD frames using the AD_Accept event. Similarly, the BC_Accept event could be used for type-BC frames.

In the current FOP-1, when FOP-1 passes a type-AD frame to the Lower Procedures it also sets the AD_Out_Flag to “Not Ready”. In this option, a count would be added at this point: when the next AD-Accept event occurs, the same frame is passed to the Lower Procedures again and the count incremented. This loop continues until the desired number of systematic transmissions is reached or until some other FOP-1 event causes a reset of the count.

The main part of this option can be specified by making changes to the FOP-1 actions “Look for FDU” (for type-AD frames) and “Look for Directive” (for type-BC frames). It will need a count variable, which will need to be set or reset in other FOP-1 actions. For example, the count needs to be reset in the action “Remove Acknowledged Frames from Sent_Queue”.

3.2.3 New FOP-1 parameter(s)

If the systematic retransmission is implemented inside FOP-1, using the method described in section 3.2.1 or 3.2.2, then there will be a number of new FOP-1 variables, and at least one of the new variables will need to be a parameter that can be set by the FOP-1 users and/or management. Therefore, there will need to be a new FOP-1 directive for setting the parameter. The DIRECTIVE.request primitive has one parameter (*Directive qualifier*) to carry the new value. The new directive will result in a new event in the FOP-1 State Table.

If multiple new variables are parameters, then there would need to be multiple new directives. It would be possible to use *Directive qualifier* to set the values of two parameters (e.g. top 4 bits and bottom 4 bits of an 8-bit value) but this is no simpler than adding two directives.

3.3 Option 3: After FOP-1

A systematic retransmission could be implemented after FOP-1, in the Transfer Sublayer or in the Synchronization and Channel Coding Sublayer.

For example, Annex C of Recommended Standard CCSDS 231.0-B-1, *TC Synchronization and Channel Coding*, defines the service interface for the services provided by the Synchronization and Channel Coding Sublayer to the user in the next layer above. The interface includes the `ChannelAccess.request` primitive:

`ChannelAccess.request` (*Frames*)

The definition of the primitive could be modified to include an optional repetition parameter:

`ChannelAccess.request` (*Frames, Repetitions*)

Depending on the value of the *Repetitions* parameter, the Synchronization and Channel Coding Sublayer would generate and transmit multiple copies of the CLTU containing the “Frames”.

This systematic retransmission option is simpler to specify, compared to the ones discussed in sections 3.2.1 and 3.2.2. It probably needs less changes to the standards documents, which is an advantage compared to those options. However, there are also some disadvantages to be considered, as discussed in sections 3.3.1 to 3.3.5. Section 3.3.6 discusses the options for setting the *Repetitions* parameter in the Transfer Sublayer.

3.3.1 Response to FOP-1 events

The option for systematic retransmission in the Synchronization and Channel Coding Sublayer would be less responsive to FOP-1 events, such as the removal of frames from the Sent_Queue. So, when a CLCW arrives which acknowledges a frame, the frame will still be retransmitted. This is inefficient, but should not happen often if the option is correctly applied for a mission with a long light delay.

3.3.2 Physical Channel or Virtual Channel

If the option for systematic retransmission is in the Synchronization and Channel Coding Sublayer, then the retransmission would be done at the level of the Physical Channel, after Virtual Channel multiplexing. For missions with multiple Virtual Channels, it could cause unexpected disturbance to the performance of the multiplexing algorithm, because all the systematic retransmissions of a frame must be completed before another Virtual Channel can access the Physical Channel. So, for example, a high-priority Virtual Channel could experience a longer delay than intended. Note that the delay would affect all types of frames, including type-BD frames.

Systematic retransmission inside FOP-1 is done at the level of the Virtual Channel, so the Virtual Channel multiplexing is unaffected. Another Virtual Channel can access the Physical Channel between the individual retransmissions in a set of systematic retransmissions for a frame.

Of course, for missions which make no significant use of Virtual Channel multiplexing, this issue is unimportant.

3.3.3 Setting FOP-1 timer

If the option for systematic retransmission is in the Synchronization and Channel Coding Sublayer, then it could be more difficult to set a correct value for the FOP-1 timer, because the delays in the system below FOP-1 could become more unpredictable. This depends on the increased difficulty in predicting delays caused by other Virtual Channels and therefore is similar to the previous item (section 3.3.2).

3.3.4 Multiple frames in a CLTU

The change to the `ChannelAccess.request` primitive looks simple, but it is more complicated when there is more than one frame in a CLTU. In general, a CLTU could carry a mixture of AD and BD frames, and it could carry frames for more that one Virtual Channel. This would need to be restricted when the “Repetitions” parameter is greater than 1, and this would need to be explained in the changes to the standards.

For simplicity, it would therefore be helpful to include a general restriction on the use of the optional repetition parameter. So, for

`ChannelAccess.request` (*Frames, Repetitions*)

when *Repetitions* is greater than 1, then *Frames* must consist of a single frame. The purpose of the systematic retransmission is to improve the probability that a frame will be accepted at the receiving end. Multiple frames increase the length of the CLTU, which in turn increases the probability that the CLTU will be rejected. Therefore, this restriction reinforces the aim of the systematic retransmissions. (See also section 4.2.1 below.)

3.3.5 Buffering the frame or CLTU

If the option for systematic retransmission is in the Synchronization and Channel Coding Sublayer, then some kind of extra buffer is needed in the sublayer, to hold the frame (or CLTU) that is being systematically retransmitted. For a TC Encoder which is implemented so that the lower levels "pull" the next frame when the link is ready, the Synchronization and Channel Coding part of the TC Encoder may be taking the frame(s) a few octets at a time and passing the data straight through. When the *Repetitions* parameter of the `ChannelAccess.request` is greater than one, the encoder also needs to save the data somewhere, and remember to use the saved data at the next transmission opportunity. The next transmission opportunity could be after quite a long delay, especially if using PLOP-1.

In older systems, FOP-1 is typically implemented in software but the Synchronization and Channel Coding part of the TC Encoder is sometimes implemented in hardware. The hardware could be awkward to upgrade for the change, especially for the addition of a new buffer to hold the frame(s) (or the CLTU) for the systematic retransmission.

In more recent systems, Synchronization and Channel Coding is more likely to be implemented in software, or perhaps in firmware / programmable FPGA. The addition of a new buffer is easier in this case.

3.3.6 Setting the *Repetitions* parameter

The `ChannelAccess.request` is part of the service interface provided by the Synchronization and Channel Coding Sublayer to the user in the next layer above, which is usually the Transfer Sublayer.

One of the procedures in the Transfer Sublayer would set the *Repetitions* parameter as appropriate, depending on the type of frame and on its own parameters. The procedure could be FOP-1 or it could be some other procedure in the Transfer Sublayer. For the purpose of updating the standards documents, it is simpler not to specify which procedure sets the parameter:

- If it is specified that the *Repetitions* parameter is set by FOP-1, then FOP-1 would also need a new directive so that users and/or management could set the optional repetition value.

- If it is specified more generally that the *Repetitions* parameter is set by the Transfer Sublayer, then it could be a managed parameter of a Virtual Channel, within the Transfer Sublayer.

4. Conclusion

4.1 Choice of option

Two of the options look better than the rest:

- The option described in section 3.2.2, for systematic retransmission in FOP-1, keeps the frame retransmission in the Transfer Sublayer, which is the layer with knowledge of the frame type and with responsibility for end-to-end delivery of the frame. The systematic retransmission can respond efficiently to a CLCW that acknowledges a frame.

- The option described in section 3.3, for systematic retransmission in the Synchronization and Channel Coding Sublayer, is simpler to specify and probably needs less changes to the standards documents. If the systematic retransmission is used as intended for a link with long light delay, the disadvantages discussed above can be handled. However, if the systematic retransmission is used on an unsuitable link, then the inefficiencies discussed in section 3.3.1 could be significant.

4.2 Other issues

4.2.1 One frame per CLTU

The Recommended Standard CCSDS 231.0-B-1, *TC Synchronization and Channel Coding*, allows multiple TC Transfer Frames to be carried in a single CLTU. In CCSDS 232.0-B-1, *TC Space Data Link Protocol*, the maximum number of frames per CLTU is a managed parameter for a Physical Channel.

At the receiving end, when a decoding failure occurs in the processing of a CLTU, the rest of the CLTU is discarded. So, if a link error causes a frame to be lost, then the same error will also cause the loss of any copies of the frame carried later in the same CLTU. It therefore seems reasonable to recommend that, whenever the systematic retransmission option is in use, then the option for multiple frames in a CLTU should not be in use. (See also section 3.4.4 above.)

Note that ECSS-E-ST-50-04, *Space data links – Telecommand protocols, synchronization and channel coding*, places a limit of one TC Transfer Frame in a CLTU.

4.2.2 Upper limit on the number of systematic retransmissions

In practice, the number of systematic retransmissions is likely to be 2, 3 or 4. Should the standards documents specify an explicit upper limit? Perhaps it would be best for the standards to state that the upper limit is a mission-specific parameter of the Physical Channel. This would require missions to make a decision, so setting their own upper limit on delays, etc. It would avoid putting an explicit value in the standards, which might place an unnecessary restriction on applications such as the scenarios described in section 1.3.2.

4.2.3 Systematic retransmission of type-BC frames

The systematic retransmissions are intended for type-AD frames. Should the automatic retransmissions also apply to type-BC frames?

A type-BC frame carries a COP-1 control command: either *Unlock* or *Set V(R)*. The commands are generated by FOP-1 following receipt of an *Initiate AD Service (with Unlock)* directive or an *Initiate AD Service (with Set V(R))* directive, and both directives require a confirming CLCW before the *Initiate* action can be successfully completed.

There are different operating procedures that can be used to initiate the AD service and different opinions about which procedures are more efficient or more secure. It therefore seems reasonable that the standards should permit the use of systematic retransmissions for type-BC frames: missions do not have to use the facility. The standards could state that the use of systematic retransmissions for type-BC frames is a mission-specific parameter of the Physical Channel or of the Virtual Channels.