esac

European Space Astronomy Centre
P.O. Box 78
28691 Villanueva de la Cañada
Madrid
Spain
T +34 91 8131 100
F +34 91 8131 139
www.esa.int

DOCUMENT

Solar Orbiter TeleMetry Corridor ICD (TMC ICD)

Prepared by	**Christopher Watson**
Reference	**SOL-SGS-ICD-0007**
Issue	**1**
Revision	**1.1**
Date of Issue	**08/05/2020**
Status	**Draft**
Document Type	**ICD**

European Space Agency
Agence spatiale européenne

Distribution

APPROVAL

Title	
Issue 1	**Revision** 1.1
Author Chris Watson	**Date** 08/05/2020

Approved by	**Date**
SOC Development Manager	
EPD PI	
EUI PI	
MAG PI	
METIS PI	
PHI PI	
RPW PI	
SOLO-HI PI	
SPICE PI	
STIX PI	
SWA PI	

CHANGE LOG

Reason for change	Issue	Revision	Date
Interpolation issue discussion. Released as a subminor revision to keep the v1_1 schema applicable	1	1.1	May 2020
Minor update following LTP planning exercise of Jan 2019	1	1	April 2019
First issue for signature	1	0	March 2017
Change to filenames (some instrument acronyms) Added alphabetically versions for ground test files Update schema version attribute for new issue	0	6	May 2016
Schema updates only Add namespace elements Corrected top-level element structure Corrected DOY_UTC_time pattern Update version attribute for new issue	0	5	April 2016
EUI comments Correct "extrapolation" -> "interpolation" Section on relationship with LTP Clarified imprecise use of "underrun" More explanation of the minimum curve meaning Added discussion of "start-up" margin in open points Added note on handling of "backup" LTP plans	0	4	Jan 2016
Update external draft	0	3	July 2015
External draft	0	2	July 2015
First internal draft	0	1	July 2015

CHANGE RECORD

Issue 1	Revision 1.1		
Reason for change	**Date**	**Pages**	**Paragraph(s)**
New "staircase" interpolation comments Typo correction - "represents"		9-11 9	2.3.1 2.2

Table of contents:

1 INTRODUCTION 5
1.1 Overview of planning cycles wrt TMC and TMC-M 6
1.2 Applicable Documents 6
1.3 Reference Documents 6
1.4 Acronyms 7
2 DESCRIPTION OF TMC CONTENT 8
2.1 What the TMC is 8
2.2 What the TMC is not 8
2.3 Interpolation between points 9
2.3.1 Issues with interpolation 9
2.4 Closure of flexibility 11
2.5 How to handle TAC windows 12
2.6 Zero of the accumulated curve 12
2.7 Missed target from the previous planning period and negative values in min curve 13
2.8 Crude picture of TM-corridor 13
2.9 Impact of RS-synoptics on TMCs 15
2.10 More on the minimum curve 15
2.11 Relationship between LTP and the corridor 16
3 DELIVERY ASPECTS 17
3.1 Delivery mechanism 17
3.2 Delivery timing and time-span 17
3.2.1 NMP and EMP 17
3.2.2 Cruise Phase 17
3.2.3 TMC-M delivery specifics 18
3.3 Filename 18
3.3.1 Version prefix and flagging of “backup” LTP plans 19
4 DETAILED CONTENT 20
4.1 Overview of parameters 20
4.2 Graphical representation 21
4.3 Schema 22
5 OPEN ISSUES 24
5.1 Measurability of TM production 24
5.2 Direct use of very high-performance comms periods 24
5.3 Possible expansion to HK and LLD 24
5.4 Possible expansion for Selective 25
5.5 Margin for the avoidance of underruns 25
5.6 Details of the TMC reissue when limits are transgressed or in contingency 27
5.7 Handling “backup” LTP plans 28

1 INTRODUCTION

This document defines the mission planning interface by which the instrument-level TM data production constraint is articulated to the Its, by the SOC, prior to their own planning which leads to IORs in MTP and STP.

This interface is called the **TMC** (TeleMetry Corridor).

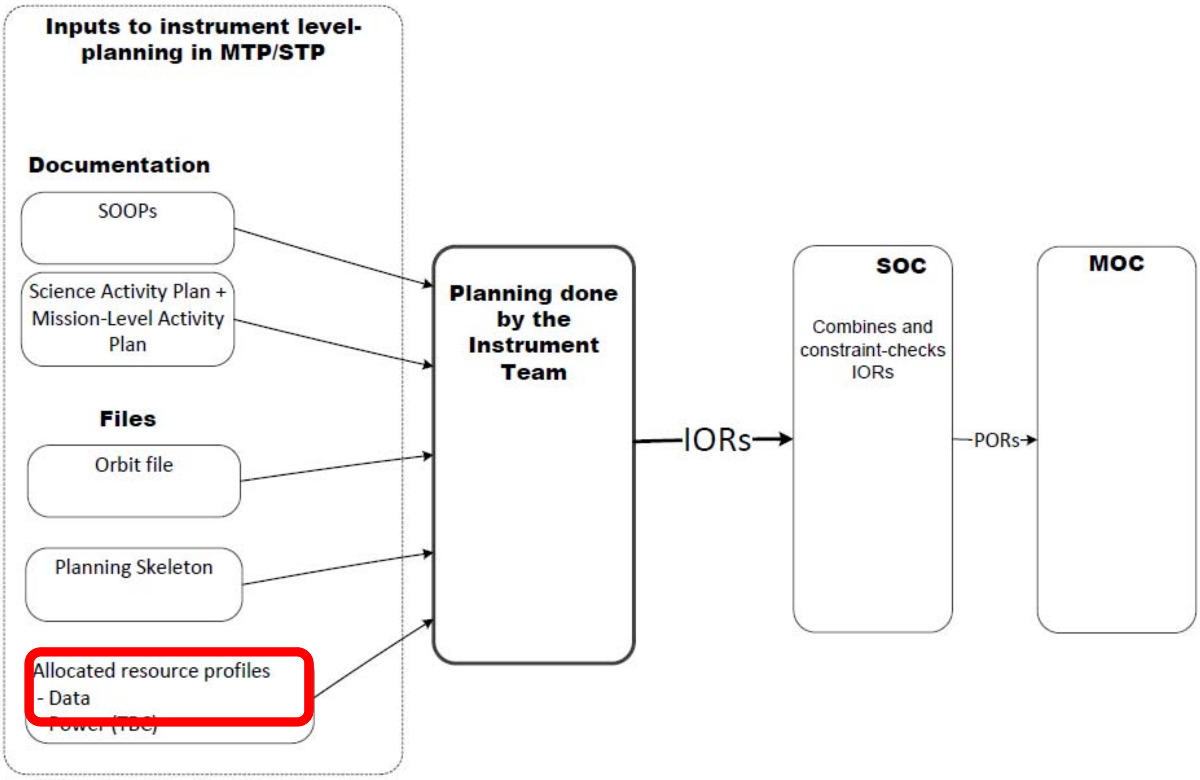

Figure 1, Instrument-team planning in MTP/STP

Figure 1 shows the context of the TMC.

The TMC is an **input**, to instrument-level planning at MTP/STP. Thus it is not foreseen to iterate the TMC to accommodate instrument teams within these cycles. Rather the work of the earlier LTP cycle is to come to a solid coordination strategy, such that such iteration is nominally not needed.

The TMC is only used in Cruise Phase, NMP and EMP (i.e. after NECP). The TMC represents the constraints over the ~six-month planning period.

Additionally a related product is generated during the onboard execution of this planning period, and shows a snapshot of the actual measured path to date (like a “you are here” marker, but including also historic “you were here” data through the planning period). This product is called the TMC-M (M for "measured"). This contains both the corridor and the measured path to date. A new TMC-M is generated following every pass, and thus is typically a daily product.

1.1 Overview of planning cycles wrt TMC and TMC-M

Planning cycle	Relationship to TMC and TMC-M
SAP / Mission-level	Minimal. *e.g. SAP sets the RSW times[1] which ultimately end up in the FECS and E-FECS*
LTP	Finalisation of **Inputs** to the TMC production
MTP	**Major use of TMC by SGS** *TMC distributed to the instrument teams at the beginning of MTP, and used to constrain instrument-team IOR planning wrt TM production.*
STP	**Major use of TMC and TMC-M by SGS** *Same TMC remains applicable to STP planning.* *Additionally the TMC-M reflecting recent measured actual TM-generation is distributed regularly.* *Exceptionally there may be an update of TMC between MTP and STP*
VSTP	Minimal. Instrument VSTP is required to be resource neutral, therefore no interaction between VSTP and TMC is expected.

1.2 Applicable Documents

1.3 Reference Documents

[GFTS] "File-Transfer SOC<->Instrument Teams ICD", Emilio Salazar, SOL-SGS-ICD-0009, v1_1, April 2018

[META] "Metadata Definition for Solar Orbiter Science", SOL-SGS-TN-0009, Sept 2018, v2.3

[1] Or more correctly, the times when the RS-instruments may write to SpW. These may not be identical to the RSWs for all RS-instruments.

1.4 Acronyms

CP	Cruise Phase
E-FECS	Enhanced Flight Events and Communication Skeleton *Also called planning skeleton. This is a SOC-extended version of the FECS that comes from MOC which details the spacecraft events*
EMP	Extended Mission Phase
HK	Housekeeping telemetry
IOR	Instrument Operational Request
IS	In-Situ
LLD	Low-Latency Data *That "thin-slice" of science data that can always be downlinked promptly to ground*
LTP	Long-Term Planning
MiB	"MibiBytes" *Non-SI unit of storage.* *1 MiB = 2^20 = 1048576 Bytes* *For info: single sector of SSMM is 1 MiB, and total available SSMM volume (2 modules) is 62 GiB*
MIB	Mission Information Database *The database defining the TM TC structures, used for example by SCOS-2000 at MOC.*
MOC	Mission Operations Centre. *For Solar Orbiter this is ESOC in Darmstadt.*
MTL	Mission TimeLine *The onboard time-tagged queue from which nominal operations execute*
MTP	Medium-Term Planning
NMP	Nominal Mission Phase
RS	Remote Sensing
RSCW	Remote-Sensing Checkout Window *The checkout windows for RS-instruments in cruise phase*
SAP	Science Activity Plan *The top-level science plan of how solar orbiter will utilise the time within CP, NMP, EMP to fulfil its science goals*
SOC	Science Operations Centre *For Solar Orbiter this is ESAC near Madrid*
SpW	SpaceWire
SSMM	Solid State Mass Memory
STP	Short-Term Planning
TAC	Turn-Around Calibration *Complements LLD as prompt science link to ground. Because a "fatter slice" of science can come through this (compared to LLD) it is normally OFF, and only enabled for specific activities that need it (those where there is a mandatory tight space->ground->space loop needed), and then only when the downlink can support it.*
TMC	TeleMetry Corridor
TMC-M	TMC-Measured *TMC plus Measured data of actual write rates seen at the SSMM*
VSTP	Very Short Term Planning

2 DESCRIPTION OF TMC CONTENT

2.1 What the TMC is

The Telemetry Corridor is a representation of the planning constraint on data production. It is designed to represent a constraint that

- Is directly controllable by the instrument (i.e. not blurred by passes or downlink-share complications)
- Is relevant to store-and-forward planning of data return managed by SOC
- Allows the individual instrument teams to plan and execute the data-production aspect of their IORs without confounding influence from whatever data-return operations other instruments are performing at the same time
- Abandons the restrictive "smooth generation" assumptions of mission-level/LTP planning to release a measure of flexibility on the TM production distribution (or broadly "over what period the average is met")

The decision therefore is to control the data production

- At the point of writing data onto SpW
- At packet-level (i.e. including packet header overheads)
- Individual corridor for each instrument
- Measured on Bulk science production only[2].
- As a corridor of allowable cumulative TM generation through the ~six month planning period defined by a maximum and minimum curve. Minimum accumulates (broadly) with the downlink allocated to the instrument's bulk store. Maximum is offset upwards from the minimum by (broadly) the size of the packet-store less a margin (to protect against overwrite).

Providing an instrument team maintains their data-production between the maximum and the minimum curves then their data-return can be considered guaranteed. If the instrument's data production moves outside of the corridor then the instrument's data-return is at risk (either because of overwrite, or because allocated downlink is not being used as expected which potentially wastes downlink). The instrument team is expected to plan operations to stay within their TM-corridor and, where necessary, take action in STP to control unexpected production trends by adjusting future generation.

The TMC is created by the SOC prior to MTP planning and is used as the data production constraint through both MTP and STP.

2.2 What the TMC is not

What the TM-corridors are *not* supposed to be representing:

[2] This does not mean that HK and LLD are "free data". It means that HK and LLD need to be subtracted from the overall allocation, before constructing the bulk-based TMC.

1. The true SSMM fill-state of the store at any moment[3]
As a consequence of having corridors that are a stable baseline for planning (i.e. that don't need to be continually re-issued), the position within the corridor (as the plan executes) will never be exactly in line with how full the SSMM packet store is in reality[4]. Rather it represents the fill-state of some imaginary store under the planning assumptions made when the corridor was created. As such true fill-state and position within the corridor will tend to correlate somewhat but only in an indicative way.
2. The exact moments when downlink occurs for each instrument
The corridor is not supposed to represent precisely how much downlink each instrument receives in each pass, nor exactly when within the pass this downlink occurs.
    a. The corridors need to contain margin to protect against the possibility of a missed pass.
    b. The SSMM fill-behaviour on Solar Orbiter is not driven by pass-to-pass storage, but rather the very long "tides" created by comms performance variation.
    c. There is no good reason to have instruments tying their SpW write times directly to passes or dump-slots within passes, when plenty of other things will be occurring in parallel .
    d. We prefer to avoid directly-coupling the instrument SpW-write planning to the SOC downlink planning. Obviously the two have to work consistently, but by retaining a small amount of margin to allow some "flex" between the two we hope for an operationally more robust system.

Consequently TM-corridors are likely to have e.g. one point per day, and not e.g. one-point per hour or one-point per minute.

2.3 Interpolation between points

Since the corridor provides max and min limits only at discrete points, it is necessary to know what rule applies in the time between the discrete points. SOC propose **linear interpolation** between the corridor points given in the TMC.

2.3.1 Issues with interpolation

The advice to use linear extrapolation is consistent with the general idea of the corridors being that:

- They reflect the long-term trending

[3] i.e. by looking at the position relative to the max and min boundaries.

[4] This is not just about the uncertainty on the overall downlink performance (missed pass etc.) It is about the accuracy with which we are able to drive the share of the downlink to respect a predefined profile. If instrument-A happens to get today a little of downlink that should have gone to instrument-B, we don't want to reissue the corridors to reflect this, only to reissue them again a few days later, once we have returned the downlink share back to where it was supposed to be.

- They are not supposed to be allowing instrument teams to aim to hit particular passes (and anyway any individual pass may fail)
- They definitively do not represent exactly when within a pass a particular downlink is going to occur.

However there are cases where linear interpolation will not work. For example Figure 2 shows an example TMC for LTP-01, with linearly-interpolated TMC boundaries in blue/red and with the generation according to LTP shown as a grey line. It can be seen that if the instrument exactly followed the LTP is would violate the low limit (when linearly interpolated).

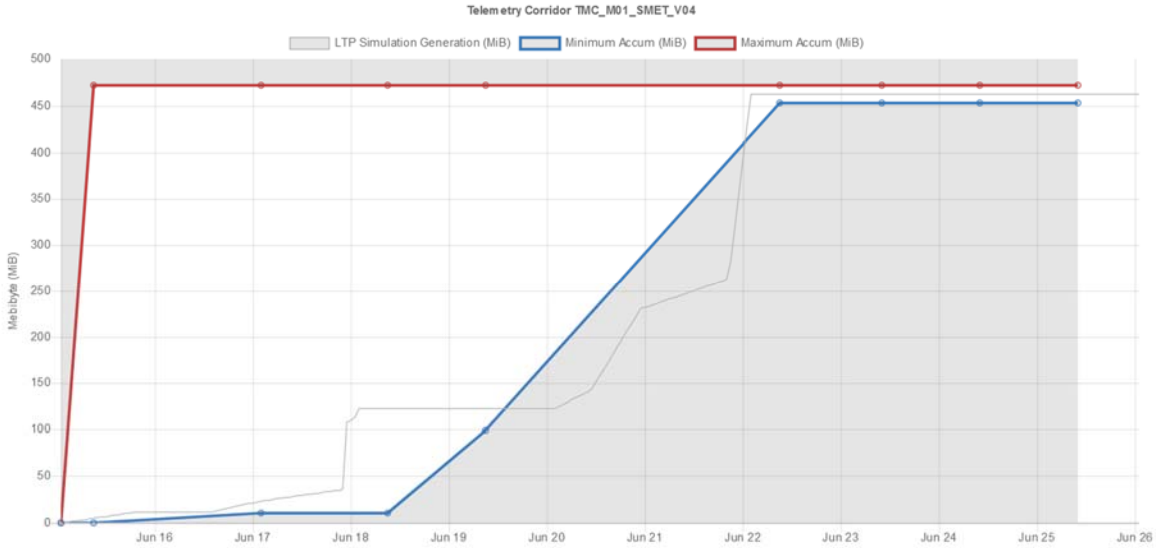

Figure 2, Example LTP-01 TMC with LTP simulation superimposed

Since the LTP simulation is the baseline plan that everything follows from, it should not occur that exactly following that plan at STP would signal a problem. Maybe the instrument team could react by generating earlier in STP to remove the apparent problem, but equally maybe they could not.

The solution to this is shown in Figure 3. Instead of linearly interpolating we implement a “staircase” approach (superimposed in orange) where the preceding value applies across intervals. I.e. each value applies unchanged, until the new value replaces it. N.b. this doesn’t affect any part of the TMC format or data content, since the data points themselves remain the same.

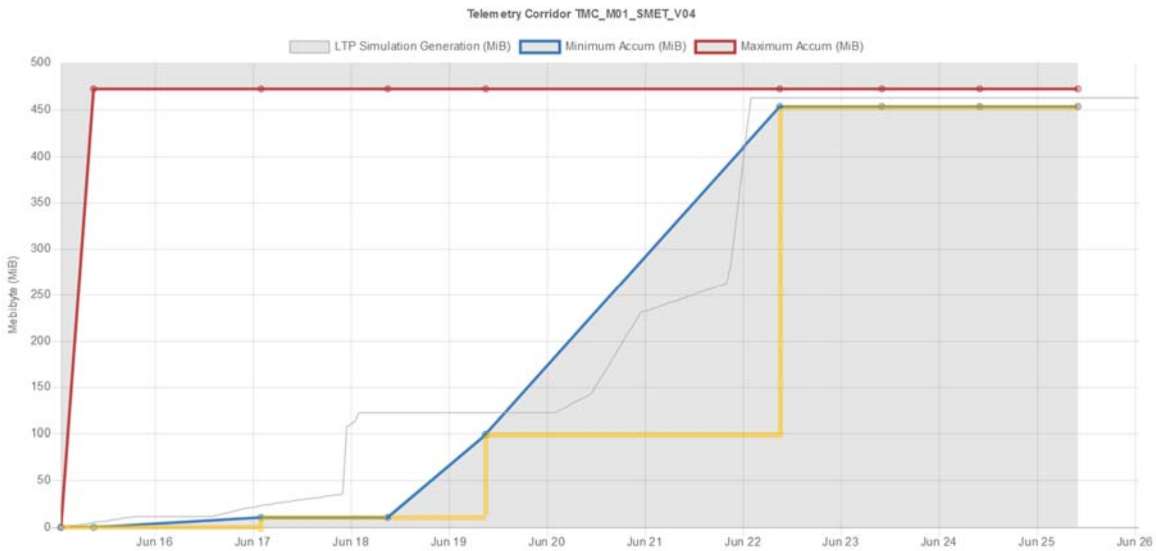

Figure 3, "Staircase" implementation of the interpolation

We expect that this solution is only ever applicable to the minimum side of the corridor.

At the current time we are not able to define firm rules about when to apply the staircase and when to apply linear-interpolation. It remains the SOC opinion that linear is better/more robust in the "benign" periods where it can be made to work.

Contributory factors that lead to this situation here are:

- Passes separated by multiple days in Cruise Phase
- Data generation in a relatively narrow interval (affects RS instruments particularly)
- Short LTP duration, combined with a LTP period that starts from empty (the problem could still manifest without this, but it makes it more obvious. There is less "space" to manoeuvre round the problem).
- Reasonable downlink performance (we are about 1 AU from the Earth here)

One could also read the open issues 5.2 and 5.5 which are closely related to this.

2.4 Closure of flexibility

There will be specific periods where the flexibility of the corridor has to be reduced to zero. Periods where this will occur include

1. At the end of each six month planning period
   We expect that each ~six month period can be planned in isolation (i.e. based on the goals of the SAP for the current period, but without impacting the SAP expectation for the next period. Because of the variations of the comms performance doesn't respect the ~six month planning boundaries, the mission-level plan will have assumed a particular fill-state at the boundary and we need to try to hit this in order not to "accidentally spend the next periods comms before it has arrived". A further concern is the uncertainty of the station-schedule for the forthcoming period.

2. For RS-instruments, between widely separated RSWs where the modelling at LTP predicts that the specific instrument's bulk-store will become empty between one RSW and the next[5]. This is necessary to preserve the principle that instruments' TM corridors operate independent of what the other instruments do. In other words, in the period where the RS-instrument store has become empty the modelling automatically assigns this downlink to other non-empty stores (as indeed the on-board systems do), and this is reflected in their downlink corridors[6]. If flexibility is allowed in these periods it means that the IS corridors are not really guaranteed.
3. During the **retained** part of deliberate SSMM-underruns. Since the stores have to be empty to allow an SSMM resizing opportunity without data loss, we have to plan to remove flexibility running into these opportunities.

2.5 How to handle TAC windows

TAC stands for Turn-around calibration. Some instruments have occasional calibration activities where

- The calibration results feed a commanding update in relatively tight-loop[7], and
- the data volume exceeds the capacity of the low-latency data

There are strict limits on how TAC is used.

For the purposes of the TMC, the TAC production is included and contributes to the overall allocation. The SOC will attempt to configure the downlink such that the occupancy of the overall downlink used by TAC, is compensated by an equivalent decrease on the TAC-using instrument's downlink used by the bulk store. This approach makes the routing of data via TAC effectively transparent to the overall cumulative situation.

Also for the TMC-M, TAC production will be added to the total production reported in the TMC-M. Providing the TAC usage is scheduled such that there is never more than one instrument writing to TAC at any time this is feasible to determine from SSMM telemetry.

2.6 Zero of the accumulated curve

We have to be clear how we define the Y-axis zero corresponding to the starting point (or left-hand side) of the corridor. On first inspection it appears easy - the corridor represents the TM generation within the period, so instruments always start a new corridor at zero.

[5] N.b. the constraint on when RS-instruments can write to SpW is not necessarily identical to the RSWs in all cases (e.g. PHI may post-process outside of an RSW). Nonetheless the principle that widely separated "write-windows" can lead to a closure of flexibility is still correct.

[6] At this point the reader may ask "can't this emptying of the RS bulk store between widely separated RSWs be avoided?" We believe that trying to achieve this will be difficult in practise, and would end up dominating the entire downlink strategy. Periods where the RS bulk stores go empty are one natural consequence of the "spiky" TM generation profiles for the RS-instruments.

[7] If the interpretation of the calibration is simple, then providing the means to apply the results autonomously onboard may be an option to avoid the delay and complication of the TAC loop to ground, and back to the spacecraft.

However, in reality there is a difference between what is known at the time of the planning, and what is known when the TMC-M with actual usage information starts to be produced. What is available to the instrument is necessarily influenced by what was done in the preceding period.

As such for planning periods where the store is nominally not empty across the boundary, the **zero** assumes that the instrument exactly hit the planned/target generation of the previous corridor. This is the same as saying that the zero reflects planning assumptions at LTP (the time the TMC is created).

> Example: We are planning the MTP-07 cycle. At the time we do this the real-life end-state of MTP-06 is not known. Therefore we assume the planned end-state of MTP-06 which we do know. Then later, to avoid unnecessarily update of products we retain the definition of the zero to be the planned boundary state.

2.7 Missed target from the previous planning period and negative values in min curve

This links to the previous point.

In the event that an instrument has missed its target from the previous planning period (and assuming that this has not led directly to overwritten data or wasted downlink), then this miss will be fed into the starting point of the measured data for the next six month planning period. More precisely if an instrument has overproduced in the previous period it may start the new period with a measured value greater than zero and vice versa.

Because SOC may allow a little flexibility around the target, and because the end of one corridor links to the start conditions of the next (plus the definition of **zero** above), this means that the minimum curve may have some **negative values near the beginning**. Of course an instrument hitting its target correctly and entering the new corridor at zero can never use this space (it cannot accumulate negatively). However an instrument that was marginally under their target (but still within the flexibility of the corridor) can meaningfully enter the new corridor at slightly negative starting value.

The extent of the actual "miss" is not knowable at the time the TMC is issued (i.e. prior to MTP), thus it makes sense to capture this carry-over within the measured data.

2.8 Crude picture of TM-corridor

Figure 4 shows a picture representing an IS TMC-M product. The planning constraint (TMC) element is shown with thick lines (Red for Max, Blue for Min). The green thin line represents the measurement element (TMC-M), the planning period being about one-third executed in this picture. It can be seen that this instrument is operating correct inside its constraint.

N.b. The figure shows relatively gentle slope variation, relative to the corridor width. Real TM-corridors may vary more strongly.

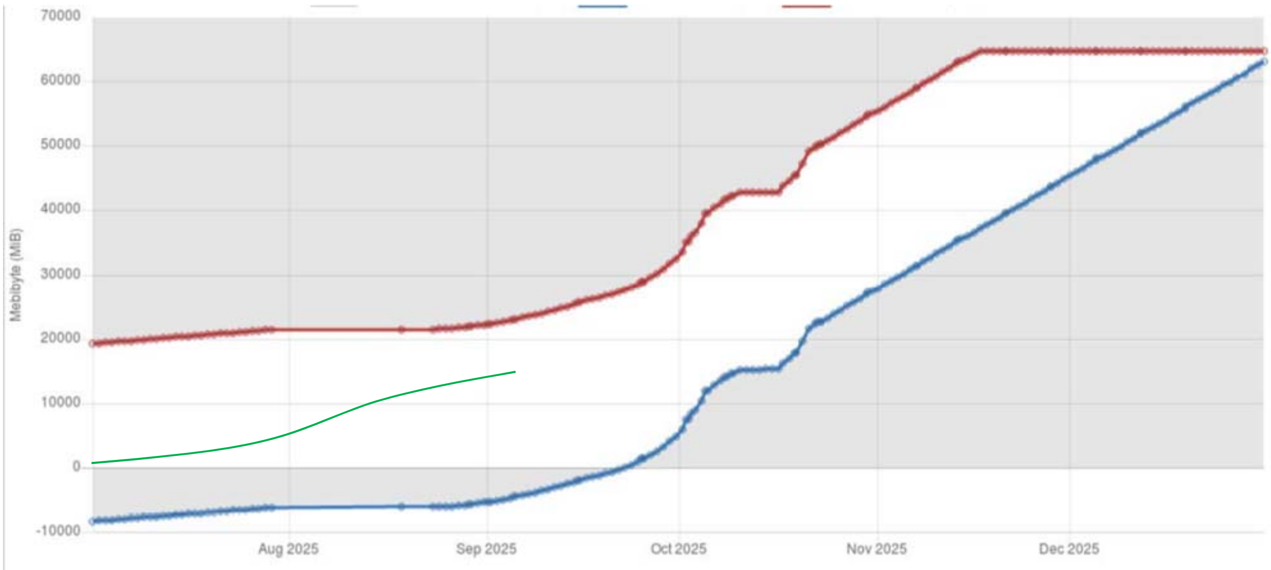

Figure 4, representation of a IS TM corridor

RS corridors, shown in Figure 5, are similar but currently we choose to "trim" the max and min curves to make them flat in the periods where science production is not allowed[8] (i.e. outside of RSWs, RSCWs, and RS_EXT extension windows, as identified in the E-FECS). This makes it easy to see the non-science periods by eye but, taking the TMC curves in isolation, it does not appear to exclude production inside of the max, min limits[9].
RS instrument teams are therefore expected to use the E-FECS and TMC **in combination**.

[8] Strictly speaking, the EFECS ICD disallows instrument commanding in these periods rather than science production. Practically speaking we believe the restrictions are equivalent.

[9] Concrete example. In the figure there is a flat period (corresponding to no RSW) centred on the beginning of October. In purely visual terms of " navigating the corridor" one could think of e.g. entering the flat period just above the min curve and continuing to produce bulk science through the flat period at a rate that avoided reaching the max curve. But this is not allowed.

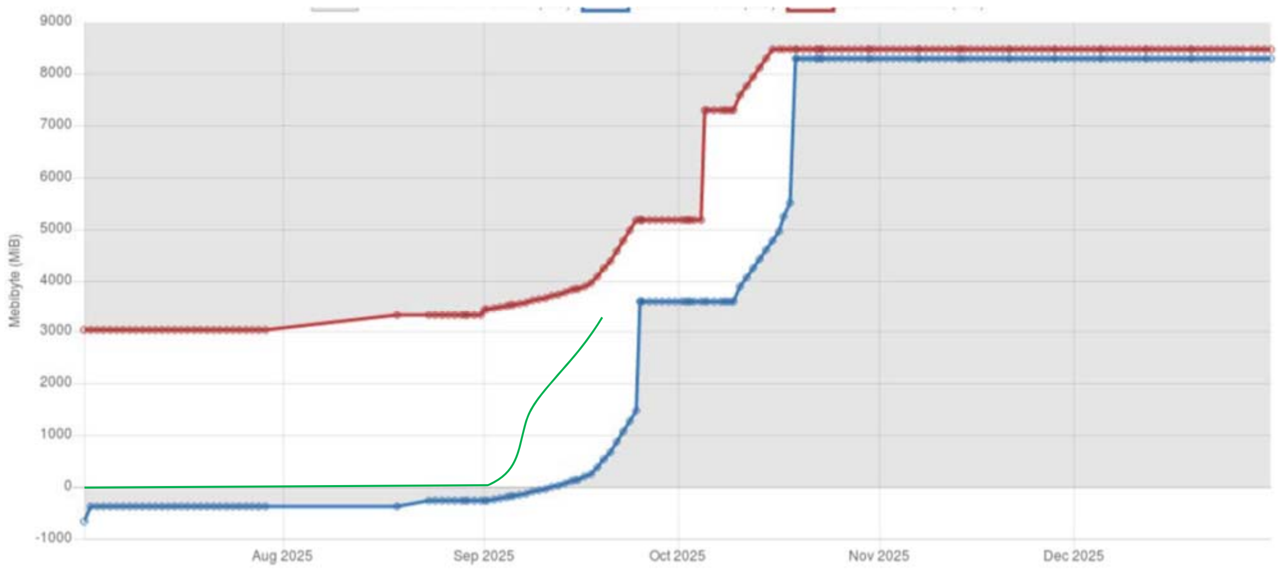

Figure 5, representation of a RS TM corridor

2.9 Impact of RS-synoptics on TMCs

Details of RS-synoptics are still in discussion. However, since the plan is to include synoptics within the existing LLD allocations, and since LLD is not included in the TMC, there is **no** direct impact of synoptics on the corridor approach foreseen.

Even if the synoptic approach were to change (placing synoptic science instead in bulk stores) probably the "flat trim" outside of RSWs etc. would be maintained. This is because the "spirit" of the synoptics is that they represent a tiny fraction of overall science production, and the current "flat trim" approximates this better than fully opening the corridor in all periods like for IS.

2.10 More on the minimum curve

The need for the maximum curve is intuitively obvious – it is a restriction to avoid that stores overwrite data.
The purpose of the minimum curve is not so immediately obvious - it is there to prevent that downlink is wasted, by there being nothing to downlink. This wasted downlink can't be recovered, and the cumulative curves of at least one TMC would have to be revised downwards.

More precisely there are two situations that the minimum curve protects against:

- That all bulk stores become empty during the pass (once HK and LL are downlinked). It may seem unlikely that all bulk stores may be empty together, but in fact there are circumstances where this is very probable. For example at the start of the nominal mission. If instruments do not start producing bulk TM at the time expected, but delay their production (i.e. fall under the minimum of their corridor) then some downlink will be wasted (since the other stores are not substantially full, having only just started their own production).

- That an individual bulk store becomes empty, when it was not expected, such that the SSMM allocates downlink to other bulk-stores. This does not immediately cause wasted downlink as in the former case. However it does mean that other instruments face a situation where their own generation constraints are no longer properly reflected by their issued corridor. This can lead to knock-on problems later. Because we want to avoid that one instrument's behaviour can impact another's published TM planning constraints the minimum curve is needed.

2.11 Relationship between LTP and the corridor

Some instruments like PHI and EUI are capable to store data internally, and then flush at a later time, potentially outside the RSW. A similar question can be raised as to, if non-baseline observing, how does this affect the corridors?
The answer is that the corridors reflect whatever is agreed at LTP- the LTP planning exercise is establishing the feasibility of plan based on the allocation of downlink capability, and then the corridor represents this downlink allocation folded together with the flexibility that can be superimposed by using the SSMM size of the bulk-store. Thus if such delayed flushes have been agreed at LTP they will be reflected in the corridor and if they have not, then they won't show up in the corridor.

3 DELIVERY ASPECTS

3.1 Delivery mechanism

It is foreseen that the file transfer mechanism is used to transfer TMCs from SOC to Instrument Teams, see [GFTS]. It is foreseen that the TMC for the specific instrument will be transferred in a package containing also the E-FECS (same file applicable to all instruments) and any other planning product applicable to MTP.

3.2 Delivery timing and time-span

3.2.1 NMP and EMP

The TMC is delivered to cover a time-range of a ~six month planning period corresponding to the station scheduling exercise. Approximately this corresponds to one period covering [Jan, June] and one period covering [July, Dec] each year[10]. Working backwards from the delivery of IORs, and using a reference time T which is the start of the onboard execution of this period:

- Preliminary MTP IORs are delivered to SOC at T-12 weeks
- Instrument teams perform their MTP planning, based on the E-FECS and the LTP plan in this interval
- TMC is delivered to the Instrument Teams at T-20 weeks

The shorter cycles (MTP and VSTP) nominally have no additional issue of a TMC. The TMC is supposed to be stable between MTP to STP (and through execution only the actual measurements are updated in the TMC-M).

Exceptionally it may be necessary to issue a revised E-FECS, updating the E-FECS that has kicked-off the MTP planning. Situations that can lead to this include:

- Mis-performance on a GAM
- Showstoppers discovered within MTP-planning, such that LTP assumptions have to be revised within the SOWG MTP planning meeting

3.2.2 Cruise Phase

Broadly the Cruise Phase mission planning works in a similar way as NMP/EMP.
However

- The top-level plan for allocation of downlink through cruise and in relation to RSCWs is probably not the SAP, but can be a document maintained by the SOC.
- A restricted/simplified LTP activity is probably still necessary every six months. This is needed to adapt plans to the confirmed station schedule. This allows to define the TMC the same as in NMP.

[10] Clearly this period is **not** the same an orbit. The above planning period has been chosen to ensure that MTP is always done with a firm station baseline already in place.

3.2.3 TMC-M delivery specifics

The TMC-M will use the same file-transfer mechanism as the TMC.
The appropriate TMC-M will be pushed to the relevant instrument team following each pass.

For information: The SOC intend to generate a PNG visually representing the data of the TMC-M for easy human comprehension. How this visual product is distributed (e.g. transferred with the TMC-M, or hosted together with the Low-latency visualisation, or both) is TBC.

3.3 Filename

Originally the intention was to adapt the Solar Orbiter metadata standard [META] to planning file naming. However [META] is not design with planning in mind, and a simpler scheme is being adopted across planning products.
The TMC filename becomes:

`TMC_M[mtp-cycle-number]_[instrum]_V[version].xml`

`[mtp-cycle-number]` is a two digit decimal number, nn, starting from 01, which increments with each new MTP[11] planning cycle.

`[instrum]` is one of

- SEPD
- SEUI
- SMAG
- SMET
- SPHI
- SRPW
- SSPI
- SSHI
- SSTX
- SSWA

`[version]` is a two character alphanumeric. Normally it is a two digit decimal number, starting from 01, which increments each time a new TMC for the cycle in question has to be issued (nominally re-issue is avoided). Versions containing alphabetical characters are used to indicate files used in ground test (avoids any possible reoccurence of a filename in-flight)

The TMC itself may be delivered in a zip file containing other planning products.

For the TMC-M the filename is:

`TMCM_M[mtp-cycle-number]_[instrum]_V[version]_[datetime].xml`

The MTP-cycle-number and the version both correspond to the underlying TMC on which the current behaviour is superimposed. There is no additional version explicitly at the level of the TMC-M, rather the additional `[datetime]` is used, on the basis that the most-recent

[11] TMC filenaming refers to MTP cycle. Other files may refer to the LTP cycle. The numbering and periods of time covered by the two are **identical** - the MTP is the "IOR filling" of the same ~six-month period as crudely planned at LTP.

measured situation is always the relevant one. The format of the datetime is yyyydddThhmmss (i.e. day-of-year date and T separator) thus an example filename is:

```
TMCM_M03_SEPD_V02_2020014T120001.xml
```

3.3.1 Version prefix and flagging of "backup" LTP plans

In particular circumstances the "V" prefix of the version may instead be a "B" for backup.

```
TMC_M[mtp-cycle-number]_[instrum]_B[version].xml
TMCM_M[mtp-cycle-number]_[instrum]_B[version]_[datetime].xml
```

Explanation:
Sometimes the science goals of a particular RSW may depend critically on the availability of a suitable solar feature that cannot be guaranteed in advance.

The proposed approach to handle these cases is to have two plans at LTP. Firstly a prime plan (with "V" files) based on the availability of the hoped-for feature, secondly a back-up plan (with "B" files) that can run reliably regardless of the actual solar conditions encountered. "B" files would only be created in the event that this potential branching of plans is agreed, and it is anticipated that the SWT would identify the RSWnthat requires the branching/existence of a backup plan before the SOWG would perform the LTP.

N.b. to prevent an explosion of permutations, only **one RSW** of a given six-month planning period would be allowed to have this dependency. The plans would be identical up to the critical RSW (and perhaps precursor), and would then diverge thereafter.

Both plans would be completed in the LTP planning, dual E-FECSs and TMCs would be delivered to the instrument teams, and dual IORs returned for the MTP planning. During the execution of the plan, it is assumed that the availability of the requisite feature can be decided by the SOOP coordinator with ~2 weeks lead time[12], such that only a single chain of STP plans have to be considered.

It is assumed that planning periods having this dependency are rare, since it implies significant additional planning work by both SOC and Instrument Teams. Also it is unlikely that we attempt this early in the mission, before we have normal planning working smoothly.

[12] This is not implying additional precursor activities. It is assumed that Earth-based assets (and perhaps Solar Orbiter low-latency data from earlier RSWs) is sufficient to decide the availability of a suitable feature.

4 DETAILED CONTENT

4.1 Overview of parameters

Header	
genTime	The time that the product was produced. Updates every time the product is created, including changes that are limited to expansion of the measurement section
ValidityRange	A YYYY-DDDTHH:MM:SS pair that identifies the (normally) ~six-month period covered. This is the span of the corridor, and the maximum span of the measured data.
MTPcycle	Duplicates the `[mtp-cycle-number]` of the filename
instrument	Duplicates the `[instrum]` of the filename
ICDVersion	Maps to the version of this document, in form vN_M
TMC (corridor) content	
count	Number of corridor triples present
creationTime	This is the time that the original TMC content was created, reflecting the definition of the corridor. Thus except in the case of reissue of the underlying corridor, this value will stay the same despite updates expanding the measurement section
*(repeat below **triple** as reflected by count)*	
x_value	Time of a corridor dataset YYYY-DDDTHH:MM:SS
min_accum_MiB	The minimum cumulative TM generation
max_accum_MiB	The maximum cumulative TM generation
TMC-M (measured) content. Present in TMC-M files but not in TMC files	
count	Number of measurement pairs present
*(repeat below **pair** as reflected by count)*	
x_value	Time of the measurement point YYYY-DDDTHH:MM:SS
measured_MiB	Measurement value

4.2 Graphical representation

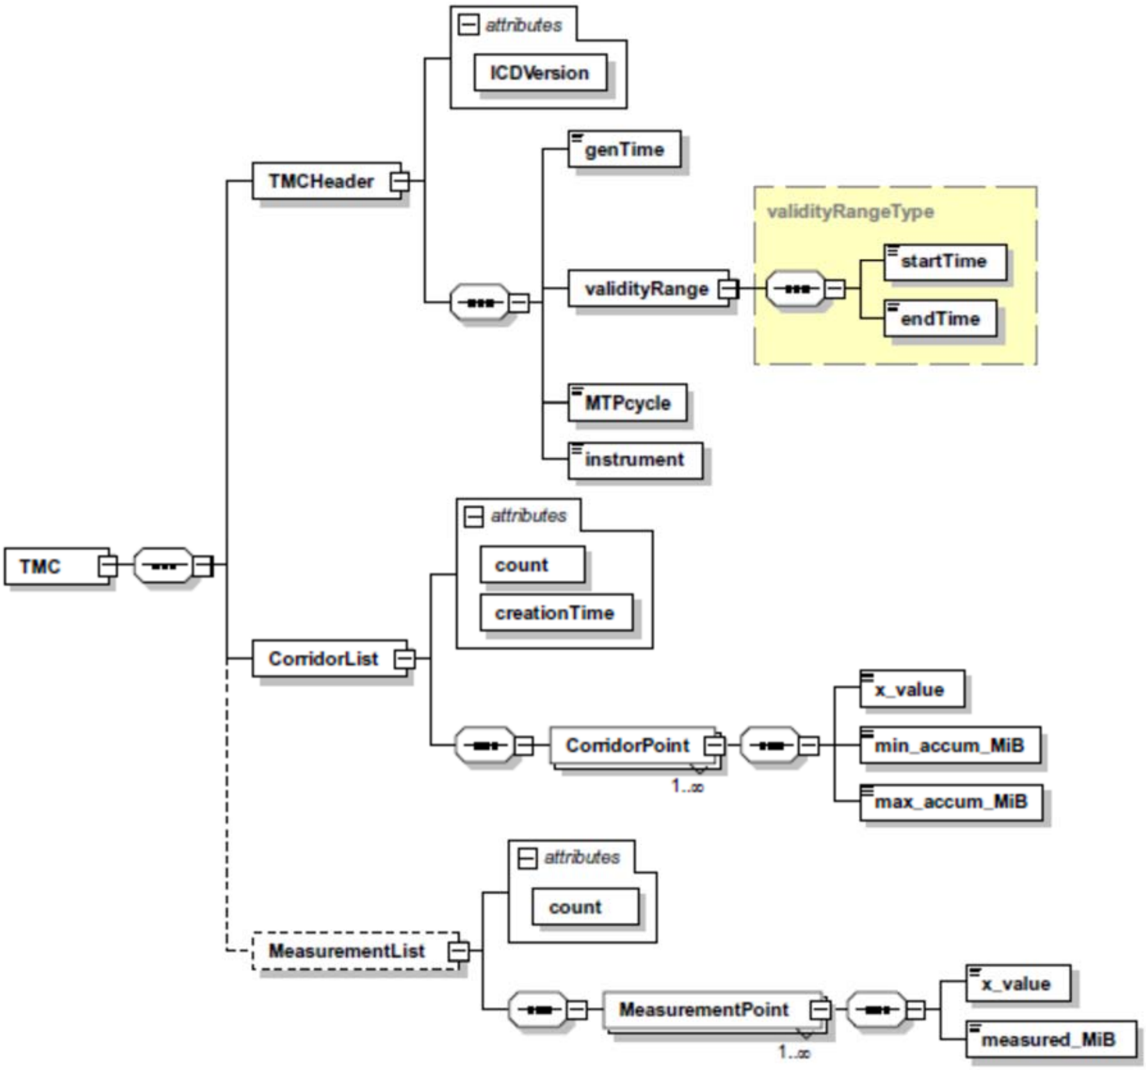

Figure 6, Schema diagram

4.3 Schema

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns="soc.solarorbiter.org"
targetNamespace="soc.solarorbiter.org" elementFormDefault="qualified">
 <xs:element name="TMC">
  <xs:complexType>
   <xs:sequence>
    <!-- Header -->
    <xs:element name="TMCHeader">
     <xs:complexType>
      <xs:sequence>
       <xs:element name="genTime" type="DOY_UTC_Time"/>
       <xs:element name="validityRange" type="validityRangeType"/>
       <xs:element name="MTPcycle" type="MTPcycleType"/>
       <xs:element name="instrument" type="instrumentType"/>
      </xs:sequence>
      <xs:attribute name="ICDVersion" type="xs:string" use="required" fixed="v1_1"/>
     </xs:complexType>
    </xs:element>
    <!-- TMC structure -->
    <xs:element name="CorridorList">
     <xs:complexType>
      <xs:sequence>
       <xs:element name="CorridorPoint" maxOccurs="unbounded">
        <xs:complexType>
         <xs:sequence>
          <xs:element name="x_value" type="DOY_UTC_Time"/>
          <xs:element name="min_accum_MiB" type="xs:float"/>
          <xs:element name="max_accum_MiB" type="xs:float"/>
         </xs:sequence>
        </xs:complexType>
       </xs:element>
      </xs:sequence>
      <xs:attribute name="count" type="xs:nonNegativeInteger" use="required"/>
      <xs:attribute name="creationTime" type="DOY_UTC_Time" use="required"/>
     </xs:complexType>
    </xs:element>
    <!-- TMC-M additional structure -->
    <xs:element name="MeasurementList" minOccurs="0">
     <xs:complexType>
      <xs:sequence>
       <xs:element name="MeasurementPoint" maxOccurs="unbounded">
        <xs:complexType>
         <xs:sequence>
          <xs:element name="x_value" type="DOY_UTC_Time"/>
          <xs:element name="measured_MiB" type="xs:float"/>
         </xs:sequence>
        </xs:complexType>
       </xs:element>
      </xs:sequence>
      <xs:attribute name="count" type="xs:nonNegativeInteger" use="required"/>
     </xs:complexType>
    </xs:element>
   </xs:sequence>
  </xs:complexType>
 </xs:element>
 <!-- Types -->
 <xs:simpleType name="DOY_UTC_Time">
  <xs:restriction base="xs:string">
   <xs:pattern value="20[0-9][0-9]-([0-2][0-9][0-9]|3[0-5][0-9]|36[0-6])T([0,1][0-9]|2[0-3]):[0-5][0-9]:([0-5][0-9])Z"/>
```

```
    <!-- the pattern accepts YYYY-DDDThh:mm:ssZ format -->
   </xs:restriction>
  </xs:simpleType>
  <xs:complexType name="validityRangeType">
   <xs:sequence>
    <xs:element name="startTime" type="DOY_UTC_Time"/>
    <xs:element name="endTime" type="DOY_UTC_Time"/>
   </xs:sequence>
  </xs:complexType>
  <xs:simpleType name="instrumentType">
   <xs:restriction base="xs:string">
    <xs:enumeration value="SEPD"/>
    <xs:enumeration value="SEUI"/>
    <xs:enumeration value="SMAG"/>
    <xs:enumeration value="SMET"/>
    <xs:enumeration value="SPHI"/>
    <xs:enumeration value="SRPW"/>
    <xs:enumeration value="SSPI"/>
    <xs:enumeration value="SSHI"/>
    <xs:enumeration value="SSTX"/>
    <xs:enumeration value="SSWA"/>
   </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="MTPcycleType">
   <xs:restriction base="xs:string">
    <xs:pattern value="[0-9a-zA-Z]{2}"/>
   </xs:restriction>
  </xs:simpleType>
</xs:schema>
```

5 OPEN ISSUES

5.1 Measurability of TM production

The TMC-M depends on new functionality being implemented in the SSMM to provide better visibility of write rates. No problem is expected, but this new functionality has to be seen in practise.

5.2 Direct use of very high-performance comms periods[13]

Section 2 describes how one-point-per-day is envisaged for TM-corridors. An exception to this rule might occur when we try to make use of these high-performance periods using TM production **specifically** in the period itself (i.e. not carrying data from an RSW to the good-comms, but generating data close to the aphelion). Because here

- the comms performance is very strong
- and the downlink could potentially be allocated to instruments having smaller stores

it is possible that the relationship between TM-generation and individual passes could become important. If this is seen to occur, then the time between TM corridor data points will be decreased for these periods.
In other words, in the normal context of the corridors the ~daily modulation of a store wrt before-pass and after-pass is negligible, compared to the longer term trends. This may break down when utilising the high performance comms periods directly.

5.3 Possible expansion to HK and LLD

HK and LLD volumes have significant effect on the overall downlink behaviour, and it is important that the SOC can model them correctly in order to properly predict bulk downlink performances (as needed for TM corridor creation). Because HK and LLD work with separate stores, are prioritised, and because they are not actively compensated against the corresponding bulk (see section 2.5) they cannot be folded into the existing corridors to present an "easy" combined TM corridor over the sum of HK, LLD, and bulk.

Some instruments have the ability to vary their HK production rates.

Broadly we expect that the rate of HK, LLD writing will be quite stable and also easily predictable already from the SOOP coordination done at LTP. Further we expect that the RS extended windows in the E-FECS implicitly define when RS-HK will be generated and when not. If this is not the enough to predict the HK, LLD rates then we may need to expand the corridor concept to present independent restrictions on HK and LLD production.

[13] These periods we sometimes refer to (imprecisely) as "underruns" in that, when we run a mission-level simulation using the EID-A rates, these periods show as underruns (or unused downlink). However in the context of later mission-planning stages where the TMC exists, the use of these periods has to have been defined in order for the capability to show up in the corridors. I.e. By this stage the underrun has been spent (if it is to be used), and the period no longer appears as an underrun. It is a job of the LTP (guided by the SAP) to make a firm allocation of these periods.

5.4 Possible expansion for Selective

Some minor enhancement of the concept is needed for the selective scheme (for those instrument that use it, so far only RPW). Most likely this involves adding an intermediate boundary/curve between the max and min cumulative curves that signals the threshold where selective is disabled/enabled. This will be added in a future version of this ICD.

5.5 Margin for the avoidance of underruns

The imposition of margin in section 2.1 is protecting against the possibility that data is overwritten within a store (so-called “overrun”).

There may be advantages to applying margin also against underruns (i.e. wasted downlink). This means raising the minimum curve above its current value. Conceptually the need is less obvious than for overruns

- Generally speaking the downlink configuration will try to ensure that if one store unexpectedly becomes empty, other non-empty stores use the downlink instead. This cannot protect in every situation (if a heavy TM producer suddenly produces nothing then the other non-empty stores may also be rapidly emptied, especially if other instruments happened to be running near the bottom of their own corridors). However there is no equivalent possibility on the “writing into stores” side.
- On the TM-production side: If an instrument does waste downlink by not producing data, then it only penalises its own downlink allocation. The easy approach is allow downlink to be wasted in this circumstance (equivalently the misbehaving instrument is immediately penalised on its allocation). The discussion here is about if and how one can do better, trying to always use downlink in the moment (by allowing other non-empty stores to use what would otherwise be wasted), and then, in later passes, correcting back to achieve the planned total allocations.
- On the downlink side: Providing the corridor has been generated on the assumption of 100% downlink performance (which effectively “biases” potential problems to the overrun side) then the underrun should only arise from minor “extra” downlink performance arising from e.g. “spare few seconds” effects like early AoS or faster than predicted OMM dump. These have no “impulsive” impact in the way that a missed pass does (although the effects may accumulate to become detectable over a six-month planning period).

However it is worth to realise that the minimum curve of the corridor is based on the smooth generation assumptions made at LTP. Thus if an instrument prefers to generate “late” with respect to what has been assumed at LTP, it is liable to violate the minimum curve, in the current approach to building the corridor. This penalty for generating late, will be more extreme in good comms periods (which are not normally our focus when worrying about downlink performance).

At the moment no margin against underrun is baselined. If such is introduced there are two obvious possibilities:

1. Base the minimum curve on e.g. 101% of downlink performance (maximum curve would stay on 100%). This means that the curves would be gently converging over time, and the corridors will narrow.
2. Offset the minimum curve upwards by some small fixed amount. This is equivalent to the approach for overruns. However it introduces a strange effect. At the start of the mission (and following every retained part of an underrun) the store is empty and the instrument begins in a situation where it is outside of allowed region. The instrument teams would then have to plan to "climb" up into allowed region "as fast as reasonable". This ramping up is analogous to the 101% convergence in the previous option, but with the instrument team deciding the best period over which to climb away from the risk of underrun, rather than the simplistic 101% applied progressively over the whole period.

Besides the above discussion of margin against underrun in general, there is a specific case (touched on in section 2.10) to do with the "start-up" from a state where a bulk store is empty (e.g. start of nominal mission, first RSW following a good comms period etc). Here the minimum curve will be based on the expected start-of-production at LTP, but maybe the instrument cannot guarantee to begin bulk production exactly according to the LTP assumption. One possibility here is to apply "start-up" margin in **time** such that the min curve of the corridor is delayed by e.g. 24 hours (or perhaps one out-of-visibility period).

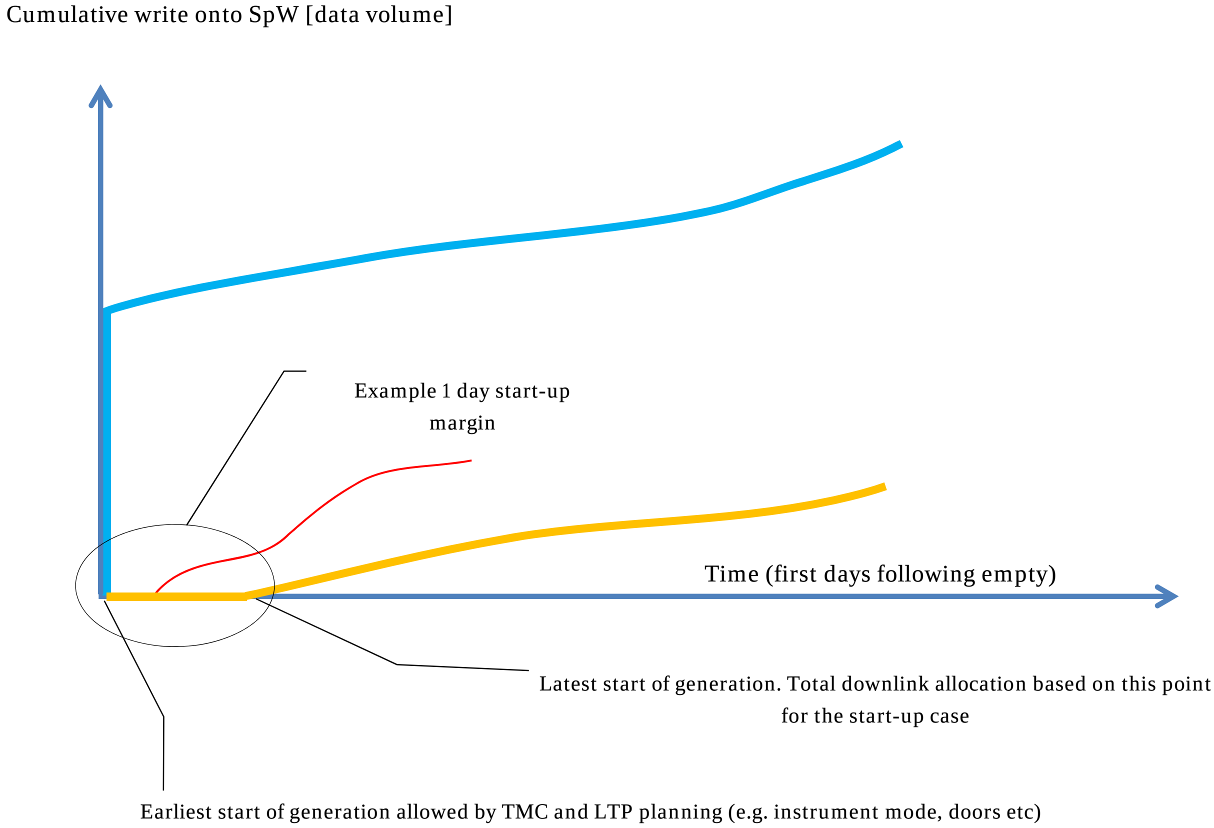

Figure 7, Representation of a "start-up" margin approach

This possible approach has to be analysed in more detail:

- The downlink modelling at LTP has to take account of latest allowed generation, to avoid that downlink usage is assumed that may later be wasted. There can be a small reduction in available TM allocation because of this approach.
- It may be possible to extend this **time** approach beyond start-up, to underrun margin in general

5.6 Details of the TMC reissue when limits are transgressed or in contingency

The principle of the TM corridors is that data-return is guaranteed providing the instrument stays within the limits of the corridor.

However the converse is not necessarily true: If an instrument deviates outside of the corridor then the data-return situation might be immediately impacted or might be recoverable, depending on how long the deviation persists, the details of what the other instruments are doing, and how rapidly the downlink can be adjusted (moving downlink in time but retaining the same overall allocations)

The current baseline is that

1. The instrument teams will see from the TMC-M promptly if they go outside of the corridor.
2. Additionally SOC will contact the instrument team to warn the that their data-return is at risk.
3. Following from either 1. or 2. the instrument team is expected to take action to bring their TM production rapidly back within bounds, utilising the next STP planning cycles[14].
4. Once SOC knows that the data-return impact is unrecoverable, they will issue a new TMC for that instrument. This TMC will reflect the data allocation loss that has occurred.

Step 4 and the process of determining whether a corridor violation is unrecoverable is a complex manual process. If it is seen to occur often in the mission, SOC may instead revise the new TMC issue to occur automatically on corridor violation.

5.7 Handling "backup" LTP plans

[RESOLVED] via "V" and "B" version prefixes.

[14] Instruments with unpredictable TM production are expected to avoid the edges of their corridors to avoid that this can occur.