

SpaceOps-2023, ID # 330

INSIGHT: THE EXPLOITATION OPERATIONS OF THE MARTIAN SEISMOMETER FROM THE SISMOC OPERATIONS CENTER

Philippe COMBES^a*, Christophe FERRIER^b,
Emilien GAUDIN^c, Frédérique MEUNIER^b, Charles YANA^b

^a Atos, for the Centre National d'Etudes Spatiales (CNES), Toulouse, 31401, France - philippe.combes@cnes.fr

^b CNES, Toulouse, 31401, France, christophe.ferrier@cnes.fr, frederique.meunier@cnes.fr, charles.yana@cnes.fr

^c Telespazio-FR, Toulouse, 31100, France, emilien.gaudin@cnes.fr

* Corresponding Author

Abstract

Launched in May 2018, the InSight (Interior Exploration Using Seismic Investigations, Geodesy, and Heat Transport) lander has deposed on the surface of Mars the SEIS (Seismic Experiment for Interior Structure) instrument after a 6.5 months travel and a challenging landing on Mars on the 26th of November 2018. NASA-JPL is leading the InSight mission, and CNES, the French space agency, after leading the development of SEIS, is now leading exploitation of measures and management of both SEIS and APSS (Auxiliary Payload Sensor Suite) instruments.

CNES has developed a complete set of tools and ground processes and integrated them in the SeIS on Mars Operations Center (SISMOC), located at the French Operation Center for Science Exploration (FOCSE), in the CNES facilities in Toulouse, France. The Ground Data System (GDS) part of the SISMOC is in charge of preparing the data for publication to the science community, especially with regards to the accurate timestamping of the instrument measurements.

Due to limited communication capabilities with Mars (bandwidth, orbiter passes), only low resolution data is received on Earth on a daily basis, and high resolution seismic data of interest can be requested after ground analysis of low resolution data. In addition, only one uplink opportunity per week is possible during the mission exploitation phase. High-resolution data requests must hence be planned carefully and data of interest transmitted back to Earth before being erased from the circular buffer on-board. The SISMOC provides science and instrument teams with a powerful tool to select the high-resolution data and build up the TeleCommands (TC) to request them.

This paper focuses on the main SISMOC functionalities with regards to the data processing and high-resolution data requests. It emphasizes how the SISMOC is at the center of a very close collaboration between the ground segment, the instrument and the science teams in order to provide efficiently the science community with high quality data focused on meaningful seismic events in the shortest possible delay.

Keywords: InSight, SEIS, SISMOC, Mars

Acronyms/Abbreviations

APSS: Auxiliary Payload Suite Sensors

E-Box: Electronics Box, the bridge between SEIS and the lander

ETHZ: Eidgenössische Technische Hochschule Zürich

FIR filters: Finite Impulse Response filters, for the conversion from an analog signal to a sampled digital signal

FRB: Full-Rate Buffer

GDS: Ground Data System

InSight: Interior Exploration using Seismic Investigations, Geodesy and Heat Transport

IPGP: Institut de Physique du Globe de Paris

SEIS: Seismic Experiment for Interior Structure

SISMOC: SeIS on Mars Operation Center

Sol: solar day, *i.e.* Martian day

TGO: Trace Gas Orbiter

UTC: Universal Time Coordinated

VPN: Virtual Private Network

WTS: Wind and Temperature Shield

1 Introduction

Interior Exploration using Seismic Investigations, Geodesy and Heat Transport (InSight) is the 12th mission of the NASA's Discovery Program [1]. It aims at improving our understanding of the formation of terrestrial planets by the unveiling the interior of Mars, its composition and structure: thickness and structure of the crust, composition and structure of the mantle, size, composition and physical state of the core.

The spacecraft is a lander with two big solar arrays and no less than 9 instruments. Most of them are on the deck. Only the Heat flow and Physical Properties Package (HP³) [2] and the Seismic Experiment for Interior Structures (SEIS) have been deployed on the ground (see Fig. 1).

In collaboration with the Jet Propulsion Laboratory (JPL), the French space agency, Centre National d'Etudes Spatiales (CNES), is responsible for two scientific instruments: SEIS and the Auxiliary Payload Suite Sensors (APSS). The latter is composed of a set of sensors that measure winds, pressure and the magnetic field, in order to shape the environmental context for the interpretation of the data from SEIS.




Fig. 2: Sectional view of SEIS on Martian ground




Fig. 1: Based on an artist's impression of the InSight lander, SEIS under the WTS and E-Box emphasized on the deck

SEIS is composed of two sets of seismometers independent of each other: 3 Very Broad Band (VBB) and 3 Short-Period (SP) sensors. It also embeds a suite of environmental sensors.

The SEIS sensor assembly is protected from the influence of the Martian weather under a Wind and Temperature Shield (WTS) – see Fig. 1 and Fig. 2.

The three VBB are enclosed in a spherical vacuum chamber. The SP are attached to the external part of that sphere. The whole sensor assembly is connected to the Electronics Box (E-Box), attached to the lander under the deck, by a tether that transmits SEIS sensors measurements. The E-Box is in charge of converting them from their analog form to a digital form suitable for the downlink to Earth.

InSight landed on Mars in *Elysium Planitia*, on November, 26th 2018. Then followed a critical operational phase of deploying the instruments on the ground. Focusing on SEIS only, it required accurate synchronization between the science, instrument and GDS teams, especially because of the operations on the sensor configuration at E-Box level. They had indeed consequences even on the GDS which required some preparation, as it will be shown further in this paper.

Since it was successfully deployed, levelled and calibrated, SEIS has recorded more than 1400 seismic events. Designed to last a Martian year, InSight extended its operational mission past that timeframe, leading to the detection of the strongest quake ever observed on another planet in May 2022 [3, 4]. SEIS has been acquiring data for more than 1440 mission Sols, until the dust on the solar array reduced the available energy so much that it prevented the lander from working.

In addition to the SEIS instrument, CNES has developed a complete operations Ground Segment. It is called SeIS on Mars Operations Center (SISMOC) and is located in Toulouse, France, as a part of the FOCSE Center, dedicated to Science and Exploration operations at CNES – see section 7.4 of [5]. On one hand, the SISMOC offers a set of basic services such as data management, task scheduling and system supervision, constituting the core system of the operation center. On the other hand, the SISMOC includes a set of mission-specific services such as management of high resolution data or correlation of the various clocks.

Even before the landing, ever since the launching on May, 5th 2018, the SISMOC has fulfilled its main two roles:

- ✓ Easing the operations through the automatizing of monitoring tasks, of the development of some kinds of TC sequences and of the delivery of files and TC to JPL for uplink.
- ✓ Processing the scientific data: sorting them and dating them as accurately as possible accordingly to their origin, their on-board processing and on-going lander and instrument activities.

Although the SISMOC has been designed to process both APSS and SEIS data and produce them in the seismic SEED format (see 3.3.1), this paper only deals with the production of SEIS data. APSS data processing is indeed more straightforward than SEIS' and only SEIS data are officially published under the miniSEED format. Through the presentation of the data walkthrough, this paper emphasizes both how the SISMOC is a key-enabler to a very close collaboration between the ground segment, the instrument and the science teams, and how it relies on this collaboration to provide constantly and efficiently the science community with data of the highest possible quality.

2 SISMOC context and overall architecture

As this paper deals with some parts of the SISMOC, it is important to give first an overall view of the platform. The SISMOC has many interfaces. It communicates with the JPL Mission Operations Ground Segment, which both provides the data received from the spacecraft and gets from the SISMOC the instrument commands to be sent to the spacecraft. It also communicates with SEIS Science entities for seismic event detection and characterization, and with APSS teams for calibration purposes, pressure and wind data formatting. It eventually distributes data to both

- the Mars SEIS Data Service (MSDS), which is in charge of producing high-level final scientific products, in order to archive them and to distribute them to the scientific community.
- the Mars Quake Service (MQS), located in the ETHZ facilities in Zürich, Switzerland, which is in charge of monitoring the SEIS and APSS data for seismic events, and of the preparation of the Mars Seismic Catalog.

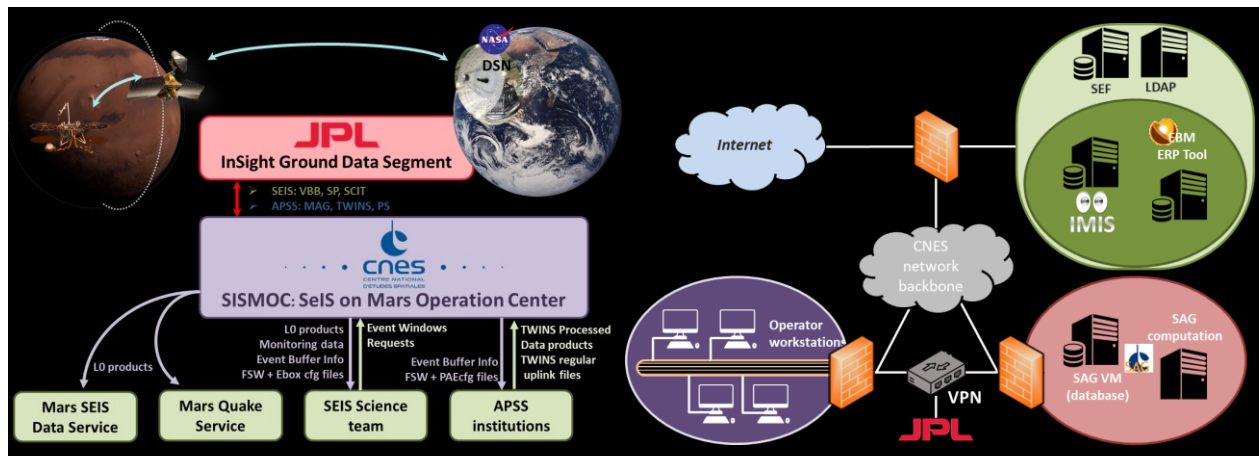


Fig. 3: SISMOC external interfaces

Fig. 4: SISMOC network architecture

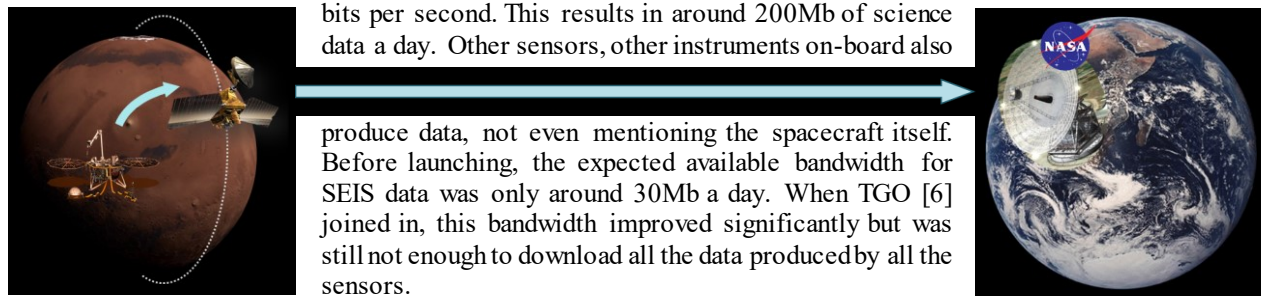
The SISMOC hardware architecture was designed to both provide a secure environment to the data and the operator workstations and still provide an external access to the published SISMOC products (on SEF server, Fig. 4), to the monitoring tool IMIS (see section 3.3.5) and to the EBM/ERP Tool server (see section 4.4.1). All the data processing is managed by the Structure d'Accueil Générique (SAG, which could be translated by Generic Hosting Structure) which runs on two different servers: one for the database and the scheduling of the processes and one for the actual computations. The former also hosts a SVN repository the use of which is explained section 4.3. Both servers are inside a protected network area with only open outward connections to the JPL data servers (through a VPN) and to the IMIS and EBM/ERP Tool servers.

3 Downlink

This section focuses on the course of data from the sensors on Mars down to Earth until they are published in a format exploitable by scientists. It is of most importance to understand how the data are collected on-board, in which conditions, how they are filtered, decimated and dated in order to understand the complexity of the ground data processing.

3.1 The Mars-to-Earth communication bandwidth, a key constraint

The spacecraft could send data directly to Earth but with such a band of frequencies that it would be only for small amounts of data. It uses Mars orbiters as relays instead. It is still not enough though to transmit all the collected data to Earth. The E-Box can be configured to acquire the data from the 3 VBB and the 3 SP at a rate of 100 samples of 24



These bandwidth constraints have impacts at various levels of the mission.

- ✓ The spacecraft hardware and software have been designed to filter out and decimate the data before sending them to Earth, while keeping on-board their full resolution version for later requests from ground.
- ✓ The GDS takes into account this spacecraft feature with regards to the computation of the real measurement timestamps.
- ✓ The operational organization includes the selection by the science team of high-resolution data to be requested for download. From the downloaded low-resolution data, the science team select so-called “events” that they would like to investigate with full-resolution data.
- ✓ The GDS embeds a software suite (see 4.4) to ease the generation and the follow-up of the TeleCommands (TC) which retrieve the full-resolution data and prepare them for download. These TC are called “event requests” in the SISMOC terminology.

3.2 On-board data processing: filtering and decimation

The E-Box is in charge of converting the SEIS sensor data from their analog form to a digital version suitable for the downlink to Earth. The digitalization chain includes of course an Analog/Digital converter, but also a set of configurable hardware FIR filters (see Fig. 5), the main role of which is the decimation of the signals down to a pre-configured rate. The filtering solely has the purpose to avoid aliasing in the decimation process.

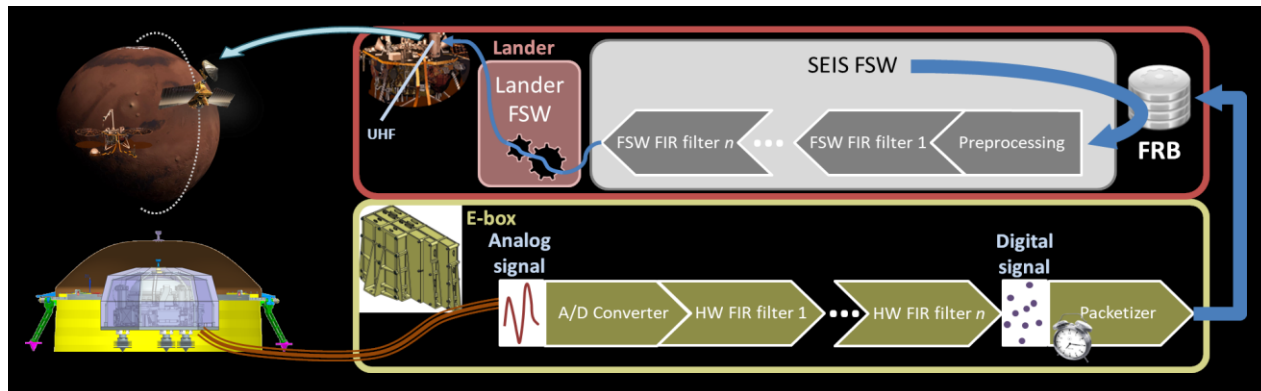


Fig. 5: On-board data processing, from SEIS to the orbiter

The data out from the FIR filters are classified in **channels**: one channel is defined by the source sensor, potentially its configuration and its target rate. Most sensors have only one possible configuration but VBB's can be set with two different gains and two different modes (scientific or engineering): a channel matches a single combination of these configuration parameters and target rates. The E-Box timestamps the data out from the FIR filters and group them by channel into packets to be sent to the SEIS Flight SoftWare (SEIS FSW). These high-resolution data packets are stored into a huge circular buffer* called the Full-Rate Buffer (FRB).

On request by the lander FSW (driven by TC), the SEIS FSW gets the high-resolution data from the FRB, drops garbage data if need be (see more details section 3.2.1), and processes the remainder through its own set of software FIR filters. In this nominal case, these filters are intended to produce low-resolution data so that it can fit in the low bandwidth of the Mars-to-Earth communications. Thus there is another set of channels defined by the outputs of the FSW FIR filters.

* When it is full, writing new data into a circular buffer overwrites the oldest data.

3.2.1 HW FIR filters delays

Since the data are only timestamped by the E-Box when they come out of the HW FIR filters, the filter pipeline shifts the signal by a so-called “**group delay**”. This delay varies with the signal source (more precisely, with the FIR filter configuration, which is set by signal source) and with the target rate (which is set through specific TC’s for the configuration of the E-Box). When converting the E-Box timestamps into UTC dates, the ground data system must be able to subtract the group delay corresponding to the signal. Although configuring the HW FIR filters is an uplink activity involving science and instrument teams (see 4.1), the subsequent group delays are required for the ground data processing.

Several events may lead to the reset of the HW FIR filters: data invalidity (sensor saturation) or discontinuity, changing the VBB gain or the target frequency, ... After these events, the HW FIR filters take some time to stabilize: this is called the “**complete FIR delay**”, which could be understood as the time spent by a sample in the pipeline. Until then, the data issued by the FIR filters are compromised. The SEIS FSW detects those events and discards in its preprocessing stage all the samples issued within the complete FIR delay after such an event occurred.

Other kinds of events may compromise the integrity of the data that the FSW cannot detect and discard likewise: calibration, VBB recentering (in order to reduce the risk that the sensor saturates), thermal compensation. As they may be of some interest for the analysis of these events, they are not discarded but flagged by the SISMOC. The SISMOC is indeed the only system in the chain that gathers all the information required to reflect these events on the produced data – see 3.3.4. But there would be no way for the SISMOC to detect those events if it was not for a hand from the instrument team.

3.2.2 E-Box status change event report

The E-Box not only processes the science signals but also generates a status report packet every time there is a status change. The status report packet contains a field that one can set through a specific FSW command: it is named *Auxiliary data* (see Fig. 7).

When setting that field to a predefined value before a calibration or a VBB recentering starts, the instrument engineer lets the SISMOC detect the event and perform the required marking of the produced data.




Fig. 6: The HW FIR filters “group delay”




Fig. 7: Byte fields of the E-Box Status Change Event Report

3.3 Ground data processing

The spacecraft data reach the JPL through the NASA’s Deep Space Network (DSN). The JPL provides then the mission partners with them. Most of these data are science data, but there are also some monitoring “House Keeping” (HK) data, such as the status change event report packet above.

The SISMOC provides the instrument team with a powerful tool to monitor all kinds of data: IMIS (see 3.3.5). The SISMOC also embeds a complete processing chain that converts science data into the widespread seismology SEED format. Fig. 8 gives an overview of this complex chain. The very first stage, the *BESTING*, is a pre-processing stage that is explained section 3.3.2. The second step, before any other computation, is the update of the Time Correlation Tool (TCT) configuration. Section 3.3.3 gives further details on this tool. This step is a key-step on the way to ensure that the timestamps are as accurate as possible in the published data. Finally, the *BESTING* output products (abbreviated TMNPF and TMPP in Fig. 8, for TeleMetry New Packets File and TeleMetry Packets Products) are sort of a reorganization of the TM packets.

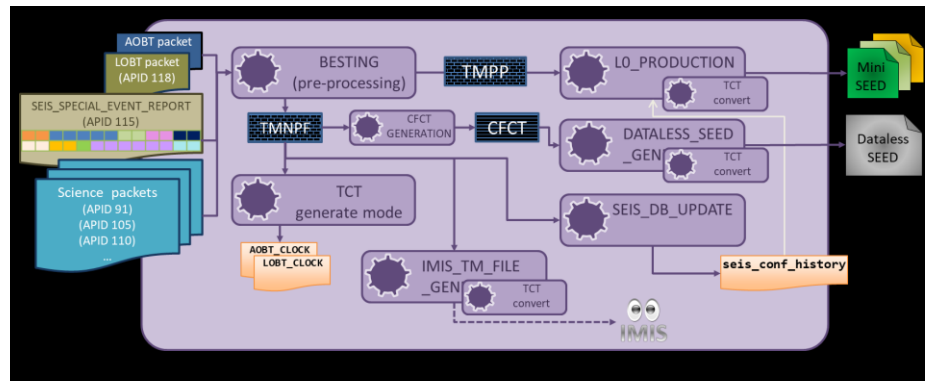


Fig. 8: Overview the SISMOC data processing chain

They are processed by the subsequent stages of the processing chain both to feed the database of the IMIS monitoring tool and to publish the data in the SEED format.

3.3.1 Overview of the miniSEED and Dataless SEED formats by the SISMOC

As it is presented in the reference manual [7], the Standard for the Exchange of Earthquake Data (SEED) is an international standard format for the exchange of digital seismological data. SEED was designed for use by the earthquake research community, primarily for the exchange between institutions of unprocessed earth motion data. It is now a format for digital data measured at one point in space and at equal intervals of time.

The miniSEED format is a subset of the SEED standard that is used for time series data. Compared to the full SEED format, a miniSEED file is stripped of most metadata that are required to interpret the data values: geographical information, response/scaling information, etc. Only minimal information such as the time series identification and some state-of-health flags are still embedded in a miniSEED file.

A SISMOC miniSEED file (see Appendix C of [5]) consists of a time-ordered concatenation of fixed length data records. A record is a continuous series of values. It has a header (see Fig. 9) with the time of the first data value in the record ("start time"), the number of samples of the series and the rate (same for all records of the file), and some reserved flags. As explained in section 3.2.1, such events as calibrations, VBB recenterings, etc. must be reported in the miniSEED file in order to avoid any misinterpretation of the data: the SISMOC uses those dedicated flags for that purpose. After the record header, the data series are compressed. The sensor, its configuration and the data sampling rate define altogether a miniSEED "channel". A miniSEED file produced by the SISMOC contains only data of the same channel (miniSEED channels are equivalent to the FSW channels – see section 3.2) on a given earth day. Hence the filename convention is *NW.STA.LOC.CHA.YYYY.DDD.r.mseed*, where:

<u>NW</u> :	Network ID	<u>STA</u> :	Station ID	<u>LOC</u> :	location ID	<u>CHA</u> :	channel ID
	2 letters		5 letters		2 digits		3 letters
<u>YYYY</u> :	data year	<u>DDD</u> :	data day of year	<u>r</u> :	file revision		
	4 digits		3 digits		>=1 digits		

There may be several revisions of the same file, either because the data have not all been received at the same time or because some new time correlation data have arrived.

A dataless SEED file is meant to hold all the SEED metadata that were removed from the miniSEED format: geographic coordinates and instrument response information. They are often needed to process reliably the time series data. A dataless does not contain any time series values but can contain a complete and comprehensive history of metadata. The SISMOC makes an extensive use of this feature.

Although there are only two sets of geographic coordinates involved (one station when on the deck and one when on the ground), the numerous instrument configuration changes or power cycles required for the operations are as many changes in the instrument response information. To reflect those in the dataless file, the SISMOC first extracts from the TM the history of the sensor configuration changes (power status, gain, acquisition rate, etc.) and computes the history of active miniSEED channels into so-called CFCT files (*CFCT_GENERATION*, Fig. 10, where CFCT stands for Configuration File Change Times). Then it combines this history with the history of the active HW FIR filters and with the SEIS FSW configuration contents to produce the dataless SEED file.




Fig. 9: Structure of a miniSEED file

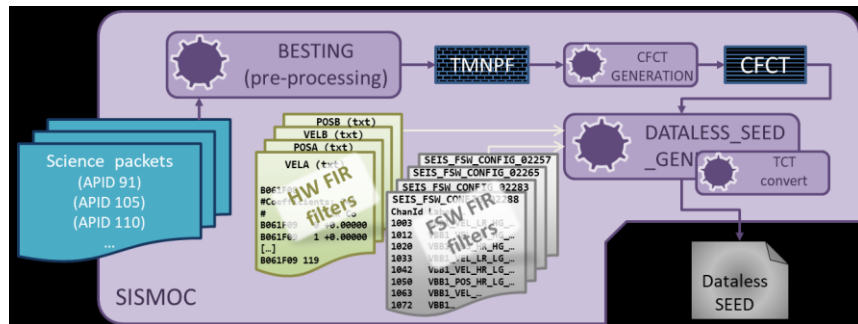


Fig. 10: Dataless generation chain

3.3.2 On-ground pre-processing

When the JPL GDS receives the TM from the spacecraft through the orbiters and the DSN, it performs only minimal checks before providing the partners with the data. The delivered files are concatenations of Space Packets that are not time-sorted and may be duplicated. Yet, every science space packet contains data from one single E-Box channel: they are all produced by the same sensor set in the same configuration and decimated at the same rate. Other space packets contain time-correlation or HK data: they are almost processed likewise.

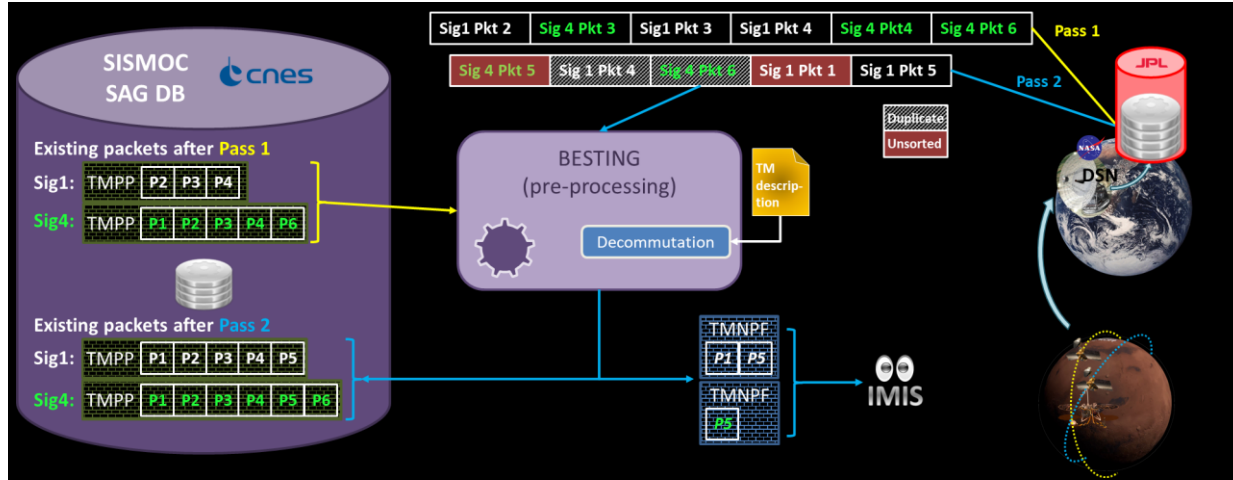


Fig. 11: SISMOC TM pre-processing

In the first stage of the on-ground TM processing (see Fig. 11), the SISMOC drops out the duplicates and organizes the data packets following the source signal, the frequency and the packet sequence number (this number was set by the packetizer on-board). The ordered packets are roughly grouped by earth day and stored in the database: yet, since it is not possible to split space packets at this stage, some data may be on the day before or the day after. The exact first and last data UTC timestamps are registered in the database to provide the subsequent processing stages with easy keys to the selection of the files to be processed. This implies that the data timestamps be converted already to UTC dates, at least the first and last of every packet.

The packets identified as new by the pre-processing, science or HK, are fully decommutated into series of couples (UTC date, value). These series are sent to the IMIS server (see 3.3.5) for monitoring purpose.

3.3.3 Date conversions

As the miniSEED files are based on UTC dates, converting the on-board timestamps is a key-point of the data processing by the SISMOC. There are three clocks involved (see Fig. 12):

- the Spacecraft CLock (SCLK) is the reference on-board,
- the E-Box clock (LOBT) is synchronized with SCLK,
- the PAE (APSS electronics) clock (AOBT) is synchronized with LOBT.

The synchronization LOBT/SCLK and AOBT/LOBT is reported by the lander FSW through specific space packets: the time correlation packets.




Fig. 12: On-board clocks and time correlation

The SISMOC embeds a tool dedicated to date conversions: Time Correlation Tool (TCT). The SISMOC retrieves the time correlation packets from the JPL server and runs TCT. The tool decommutates the packets and updates its internal APSS and SEIS time correlation tables. These tables are the main input to any conversion from AOBT or LOBT dates to SCLK dates. A third table lists the correlation points between the SCLK time and the SpaceCraft Event Time (SCET) with additional information that make the conversion from SCLK to UTC straightforward. It is provided by the JPL team (see Fig. 13) on a 3-month period basis. That period is actually variable because the decision to publish a new version of this correlation file is driven by the drift between the SCLK and SCET crossing a threshold value. Every time a new version is issued, some data must be reprocessed in order to improve the accuracy of their UTC timestamps. Although the SISMOC has been designed to identify automatically the data that should be reprocessed,

the delivery period length is such that too many files are involved: the SISMOC reaches some internal safeguards and a manual reprocessing by the GDS team is required.

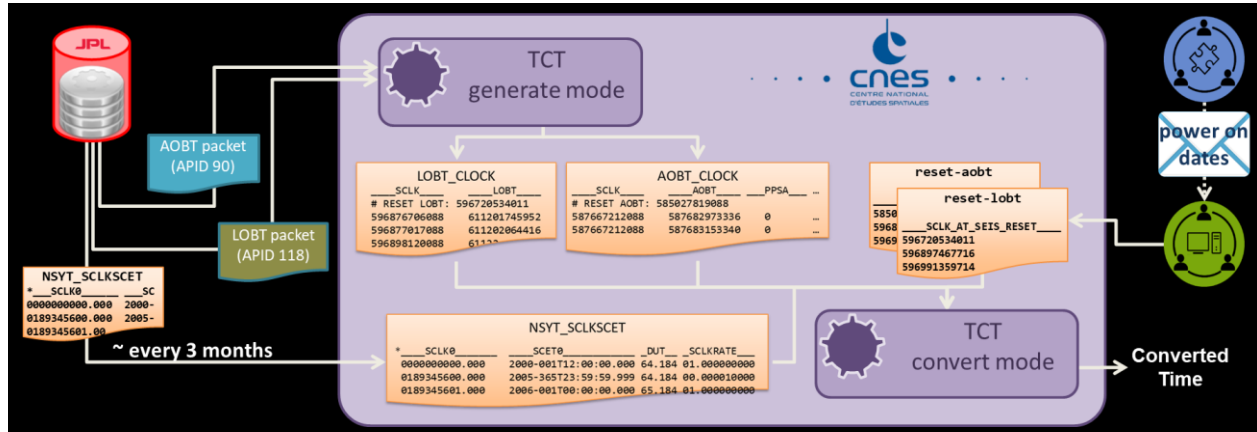


Fig. 13: Time Correlation Tool guidelines

A conversion based on correlation tables relies on usual interpolation mechanisms, such as cubic splines, when the date to be converted is between some points of the table. It relies on extrapolation mechanisms, such as the least squares method, when it is after the last point of the table (see Fig. 14). There are some cases however when these basic principles would lead to inaccurate results. Every time the E-Box is powered on after a while, the drift between LOBT and SCLK has to be reset. The operational team know of the exact power on time and inform the GDS team so that they can update the TCT configuration to better the correlation between both clocks. Note that Fig. 14 illustrates this process for the E-Box and LOBT only, but it is the same for the PAE and AOBT.




Fig. 14: TCT date conversion algorithms

3.3.4 MiniSEED generation

All dates in SEIS miniSEED files are UTC dates converted from LOBT dates by TCT. Yet a LOBT date is not the exact date of the measurement, because the timestamp is set once the data come out from the HW FIR filters (see section 3.2.1). Therefore, the miniSEED generation process, *L0_PRODUCTION*, has to subtract the group delay from the SEIS dates. From the TM it gets which FSW channel the data belong to and the active FSW configuration. From the FSW configuration files stored in the SISMOC DB, it retrieves the relevant group delay. Once again, it is noticeable that the dating process relies on a strong collaboration between the teams: these FSW configuration files are generated by the instrument team, based on information and values provided by the science team (see section 4.2).

As section 3.3.1 explains it, miniSEED records have a header containing the start time, the number of samples and their rate. Therefore, there are two ways of computing the date of the N^{th} sample:

$$UTC_N = UTC_{start} + \frac{N-1}{Rate} \quad \Bigg| \quad UTC'_N = tct_convert(LOBT_N)$$

Whenever the drift between UTC_N and UTC'_N crosses the threshold of 1%, the SISMOC truncates the record, pads it with zeroes up to its fixed length and starts a new record from the N^{th} sample with $UTC_{start} = UTC'_N$.

The *L0_PRODUCTION* may also truncate records because of on-board events. The miniSEED record headers also have flags used to report those on-board events that may have had an impact on the measurements. Some space packets are dedicated to carrying this information: the E-Box status change event report packets (see 3.2.2). Once they have been pre-processed by the *BESTING* stage, another SISMOC process, *SEIS_DB_UPDATE*, decommutates those packets to maintain the list of all these events (see the file *seis_conf_history* in Fig. 15). The *L0_PRODUCTION*, uses this list to identify which records are impacted by events among those that it is currently processing.

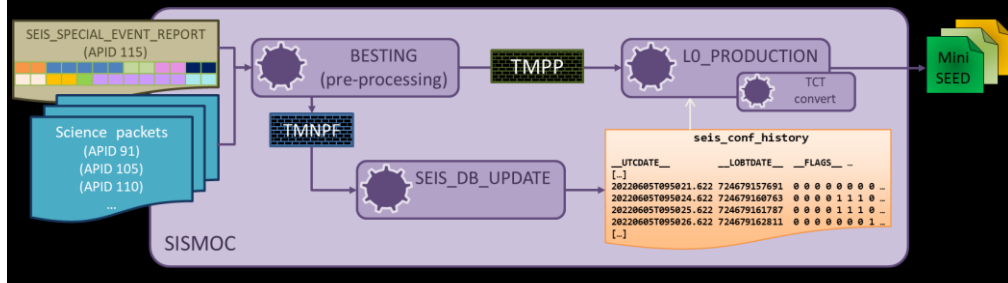


Fig. 15: MiniSEED generation chain

The first step here is to compute the exact time range when the data are impacted by an event. Section 3.2.1 presents the complete FIR delay as the stabilization time of the HW FIR filters after an event. That delay must be added to the event end time so that the data being compromised by the stabilization of the filters are flagged too. Again, the *L0_PRODUCTION* retrieves the complete FIR delay from the FSW configuration files, a teamwork product.

The second step is to flag the right records. A naive approach would flag all the records that overlap the computed event time range. But the time range may start and/or end in the middle of a record. If such a record were to be flagged as is, then all the samples of the record would be considered as impacted by the event (see Fig. 16). In order to flag only those samples that were actually impacted by the event, the *L0_PRODUCTION* truncates the records at the event start and end dates. This produces incomplete records, padded with zeroes up to their fixed length.

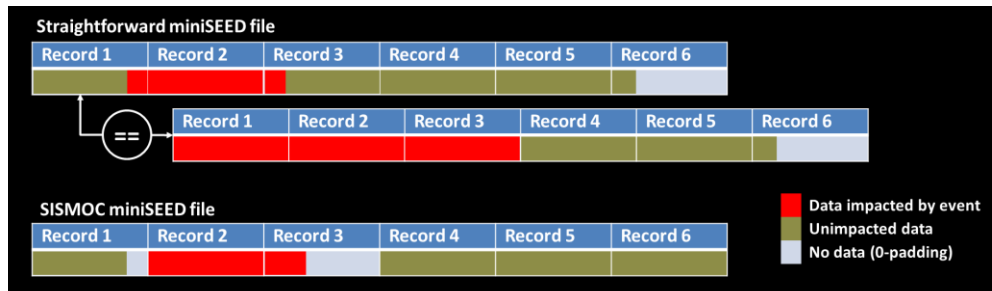


Fig. 16: MiniSEED record splitting on event dates

In this section and the previous one, we have focused on the generation of the scientific SISMOC products, how the SISMOC ensures their accuracy and thus the quality of their exploitation by the science community. These are the most important products of the SISMOC, meant to last long after the mission is over. During the mission however, the SISMOC produces all sorts of products destined to be used for the operations only, such as those that feed the main monitoring tool: IMIS.

3.3.5 Monitoring with IMIS

The SISMOC sends all TM data to the Instrument Monitoring Interactive Software. This multi-mission tool had already been adapted to the specificity of the Martian missions with Mars Science Laboratory (MSL). It is based on a classical client/server architecture. The client is an Eclipse based application. It offers various visualization possibilities,

such as plotting specific telemetry data with advanced chart functionalities, listing values chronologically, defining templates that gather some views to focus on specific aspects of the instrument or measurements, etc. The instrument team make an extensive use of the template feature (see Fig. 17) because IMIS can generate reports from them.

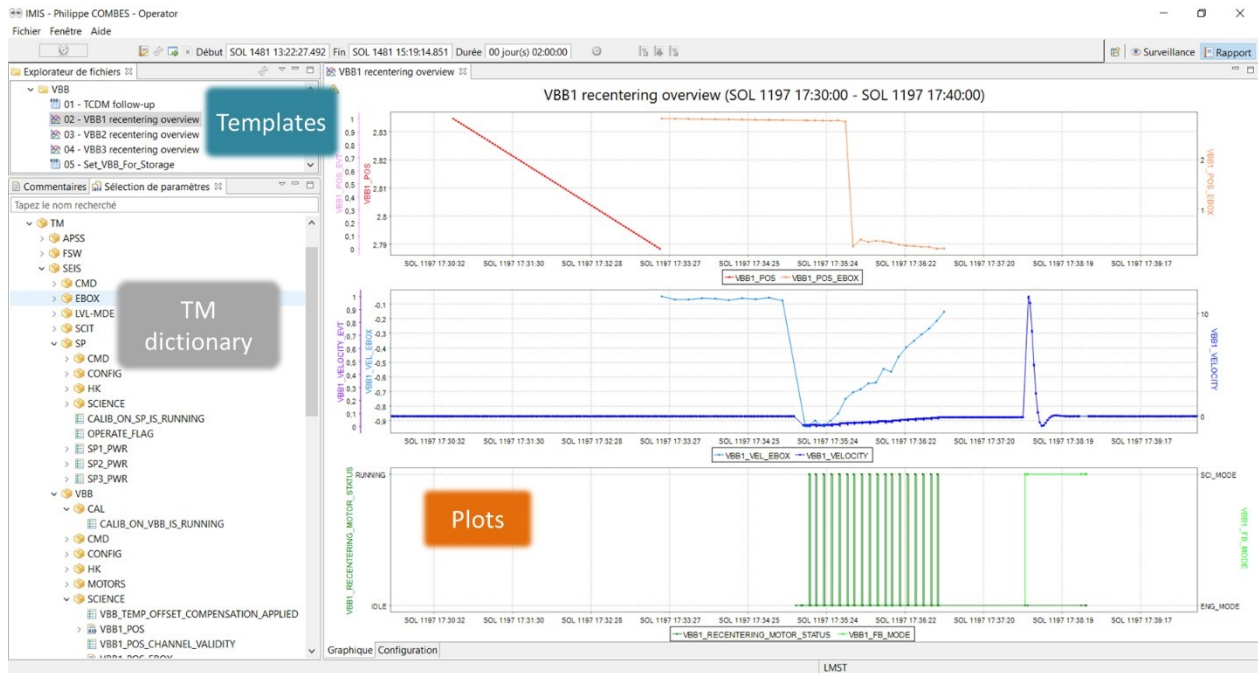


Fig. 17: IMIS snapshot: instantiation of a template with several plots

With the client, the user can also define monitoring rules to raise alarms on TM reception. The server gets the data as couples (date, value) from the SISMOC process *IMIS_TM_FILE_GENERATION* (see Fig. 8), then sorts them following the classification of the TM that was decided at installation time (see the TM dictionary on Fig. 17 and Fig 18). It applies on the fly the defined monitoring rules and raises alarms consequently. It can even be configured so that an e-mail is sent whenever a new alarm has been raised.

With all these features, IMIS lets the instrument team monitor accurately the health status of the instruments as well as the on-board execution of the command sequences. A command line version of the IMIS client is used on a weekly

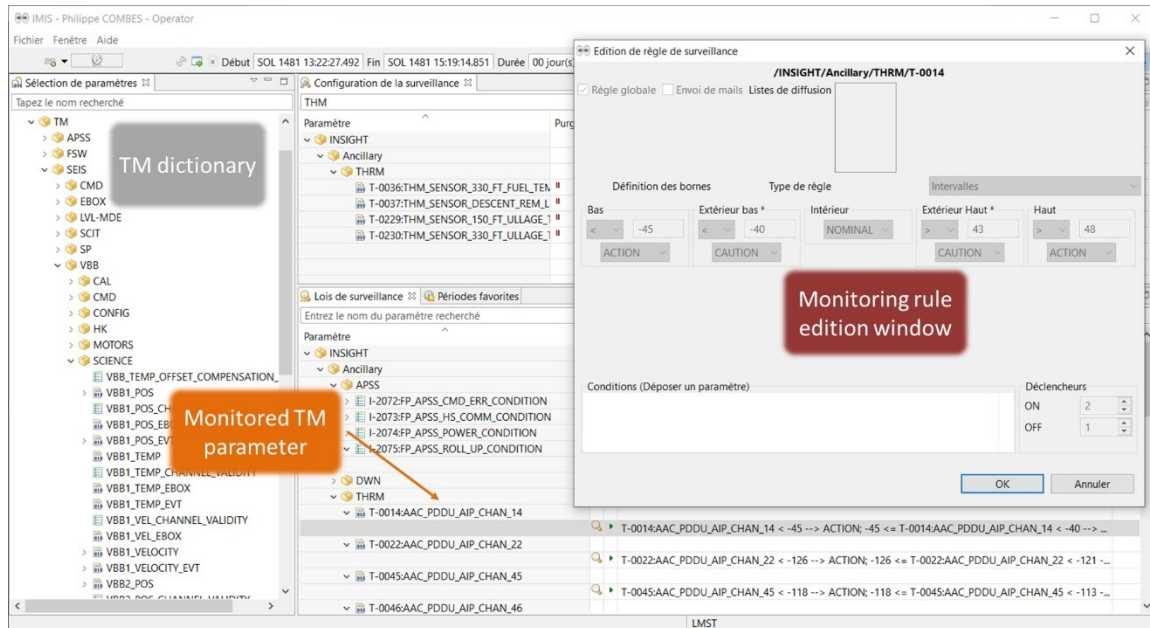


Fig 18: IMIS snapshot: example of monitoring rule definition

basis to generate a full status report with plots, a summary of alarms, etc., thanks to the templates prepared with the graphical client. It also offers a report mode that is used every year, when a full report on the instruments is needed to monitor the evolution over long periods of time of the consumables or the ageing of the instruments. The command line client is not only useful to the instrument team. The SISMOC relies on it to generate the lander activity files.

3.3.6 Lander activity files

If the miniSEED files and the dataless SEED are the main products of the SISMOC, it also publishes various kinds of files that may be useful for the operations or for the analysis of the seismic data.

A lander activity file is simply a list of predefined activities with their start and end dates. They are published under the QuakeML format [8], which is a subset of XML. These files are generated once a day. The generation process first computes the oldest and newest dates of all TM received since its previous execution. Then it relies on the IMIS command line client to request for all activity starts and ends that occurred between these two dates. The key-point here is the identification of these activity starts and ends. Fortunately, this sensitive task was to be done only once, but it is noticeable that all teams committed themselves into it.

First, the science team defined a list of lander activities they were interested in. The instrument team then identified the TM parameters and the conditions on these parameters that would reflect the start or the end of an activity. The team then used the IMIS graphical client to define monitoring rules based on these parameters and conditions. With this method, the alarms raised on TM arrival mark actually the expected activity start and end events. It is rather diverting the IMIS alarm system from its original purpose because these alarms have nothing to do with the operational instrument monitoring. This is why they are only visible to a specific GDS application user, the one used by the generation process when it sends its requests. These requests work with IMIS templates that shape the results in a format suitable for the parsing by the generation process. With the help of the GDS team, the instrument team used the IMIS graphical client to prepare those templates.

Every day since that unique configuration step, the GDS team has been also in charge of supervising the correct execution of the lander activity files generation process.

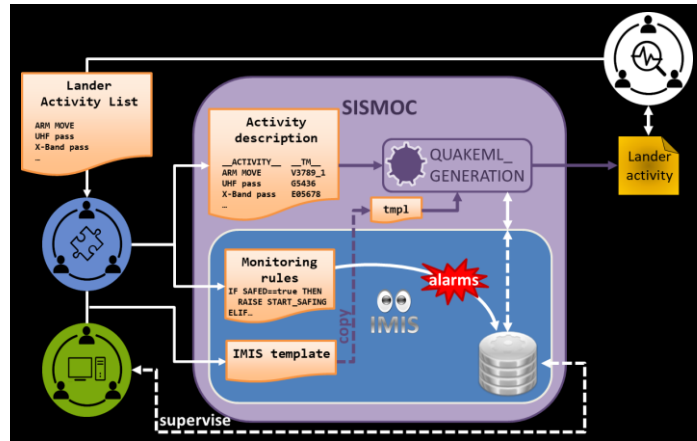


Fig. 19: Lander activity workflow

4 Uplink

Several uplink activities rely entirely on the SISMOC environment and tools and they may have major consequences on the downlink processing chain as it was emphasized all along section 3.

4.1 Configuration of the HW FIR filters

The configuration of the HW FIR filters is yet another good example of the collaboration between all committed teams. It has been necessary to update it several times after the deployment of the instruments, especially when HP³ was hammering its mole into the ground. The science team define the coefficients. The instrument team use the SISMOC tools to build the binary files from those coefficient lists, to store them in the

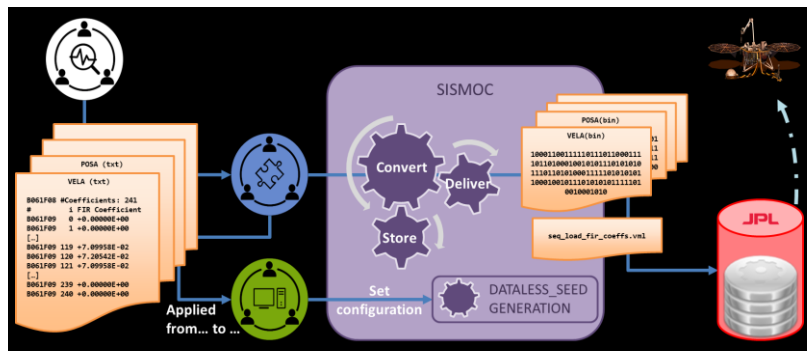


Fig. 20: HW FIR filters configuration

version management system, to build the TC sequence that shall apply the new configuration files to the HW FIR filters and finally to deliver the files and the sequence to JPL for uplink. Meanwhile, they also are responsible for informing the GDS team of the dates when the filters shall be applied on-board. With this information joint to the FIR filter contents, the GDS team are able to update the configuration of one of the SISMOC processes: the dataless generator (see 3.3.1).

4.2 Configuration of the FSW FIR filters

The configuration of the FSW FIR filters is a bit more tedious. It is actually part of the configuration of the SEIS FSW as a whole. The SEIS FSW can process the E-Box and PAE channels in various ways, not only through its FIR filters. In practice however, the channels are either not processed at all (kept “raw”, as they come out from the E-Box or the PAE) or processed through the FSW FIR filters. A FSW configuration file is therefore mainly a configuration file for the FSW FIR filters.

Such a file is the binary version of a table (see Fig. 21) that specifies which channels are to be produced. For every output channel, the table gives the source E-Box or PAE channel(s), all the processing stages they shall go through and their parameters if any. The table also provides the group delay and the complete FIR delay of every channel. Those delays have been explained section 3.2.1 for the HW FIR filters. The FSW FIR filters have the same kind of delays. The table sums up the delays for both HW and FSW FIR filters. This is why it so important for the SISMOC to have access to these files when processing the TM: as explained in section 3.3.4, group and complete FIR delays

output channel ID		output rate	complete FIR delay	group delay		decimation ratio		output channel name	input channel name & ID	
channel id	frequency	complete FIRDelay	groupDelay	algorithm Type	filter	averaging WindowSize	downsampling Ratio	label	entryChannels	input ChannelId
1002	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_LR_HG_SC_10Hz	VBB_1_VEL_20_HZ_SCI_HIGH_RAW	
1013	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_HR_HG_SC_10Hz	VBB_1_VEL_100_HZ_SCI_HIGH_RAW	1012
1020	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB1_POS_HR_HG_SC_pt5Hz	VBB_1_POS_1_HZ_SCI_HIGH_RAW	2
1032	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_LR_LG_SC_10Hz	VBB_1_VEL_20_HZ_SCI_LOW_RAW	96
1043	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_HR_LG_SC_10Hz	VBB_1_VEL_100_HZ_SCI_LOW_RAW	1042
1050	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB1_POS_HR_LG_SC_pt5Hz	VBB_1_POS_1_HZ_SCI_LOW_RAW	97
1062	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_LR_HG_EN_10Hz	VBB_1_VEL_20_HZ_ENG_HIGH_RAW	108
1073	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_HR_HG_EN_10Hz	VBB_1_VEL_100_HZ_ENG_HIGH_RAW	1072
1080	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB1_POS_HR_HG_EN_pt5Hz	VBB_1_POS_1_HZ_ENG_HIGH_RAW	109
1092	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_LR_LG_EN_10Hz	VBB_1_VEL_20_HZ_ENG_LOW_RAW	126
1103	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_VEL_HR_LG_EN_10Hz	VBB_1_VEL_100_HZ_ENG_LOW_RAW	1102
1110	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB1_POS_HR_LG_EN_pt5Hz	VBB_1_POS_1_HZ_ENG_LOW_RAW	127
1120	0,1	0	50	NO_PROCESSING				1 VBB1_TMP_LR_pt1Hz	VBB_1_TEMP_PT1_HZ_RAW	13
1130	0,1	423000	210050	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB1_TMP_HR_pt1Hz	VBB_1_TEMP_1_HZ_RAW	1131
2002	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_LR_HG_SC_10Hz	VBB_2_VEL_20_HZ_SCI_HIGH_RAW	4
2013	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_HR_HG_SC_10Hz	VBB_2_VEL_100_HZ_SCI_HIGH_RAW	2012
2020	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB2_POS_HR_HG_SC_pt5Hz	VBB_2_POS_1_HZ_SCI_HIGH_RAW	5
2032	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_LR_LG_SC_10Hz	VBB_2_VEL_20_HZ_SCI_LOW_RAW	98
2043	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_HR_LG_SC_10Hz	VBB_2_VEL_100_HZ_SCI_LOW_RAW	2042
2050	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB2_POS_HR_LG_SC_pt5Hz	VBB_2_POS_1_HZ_SCI_LOW_RAW	99
2062	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_LR_HG_EN_10Hz	VBB_2_VEL_20_HZ_ENG_HIGH_RAW	110
2073	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_HR_HG_EN_10Hz	VBB_2_VEL_100_HZ_ENG_HIGH_RAW	2072
2080	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB2_POS_HR_HG_EN_pt5Hz	VBB_2_POS_1_HZ_ENG_HIGH_RAW	111
2092	10	4471	2195	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_LR_LG_EN_10Hz	VBB_2_VEL_20_HZ_ENG_LOW_RAW	128
2103	10	4471	2215	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_VEL_HR_LG_EN_10Hz	VBB_2_VEL_100_HZ_ENG_LOW_RAW	2102
2110	0,5	102800	51000	FIR	d:/seis/FIR_AA_DIV2_GA04_RFG_01.bin		0	2 VBB2_POS_HR_LG_EN_pt5Hz	VBB_2_POS_1_HZ_ENG_LOW_RAW	129
2120	0,1	0	50	NO_PROCESSING				1 VBB2_TMP_LR_pt1Hz	VBB_2_TEMP_PT1_HZ_RAW	15
2130	0,1	423000	210050	FIR	d:/seis/FIR_AA_DIV2_PFO5_RFG_01.bin		0	2 VBB2_TMP_HR_pt1Hz	VBB_2_TEMP_1_HZ_RAW	2131

Fig. 21: Table representation of a FSW configuration file

are required in the generation process of the SEIS miniSEED files.

Thus, although defining those configuration files is a task for the instrument team in collaboration with the science team (for the definition of the relevant output channels), the produced files are configuration files for the GDS as well as they are for the SEIS FSW on-board. What is more, the SISMOC provides the instrument team with a convenient graphical tool to generate these files. From a graphical representation of the processing stages, it can compute some of the data in the table, which reduces the risk of mistakes in the computations.

4.3 Commanding

There are two kinds of command sequences to be run on-board: tactical sequences and event sequences. A tactical sequence performs specific actions on the instrument; it is planned and decided in collaboration with the science team. An event sequence only consists of retrieving high-resolution data from the FRB and preparing them for downlink. The SISMOC embeds a SubVersionN (SVN) repository [9] for the versioning of all command sequences. It additionally stores their final version in its main database, a necessary step for the automatized delivery-to-JPL process.

For the development of tactical sequences, the SISMOC integrates a JPL tool, MPS Editor, that offers the instrument team a full development environment: sequence edition, checking and compiling through tools hosted on JPL servers and a direct connection to the SVN repository for the version management – see Fig. 22.

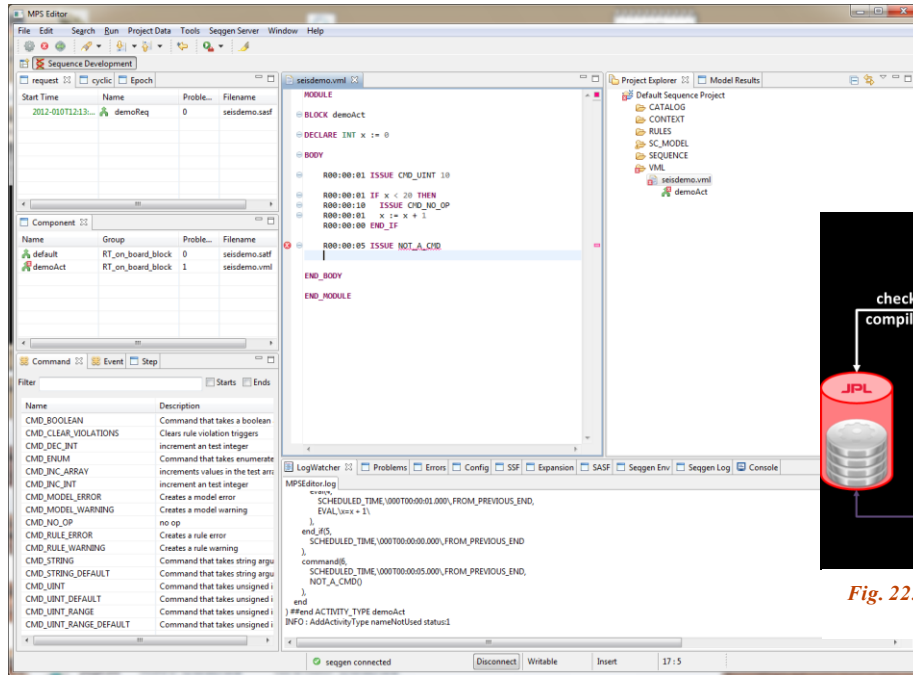


Fig. 23: MPS Editor snapshot

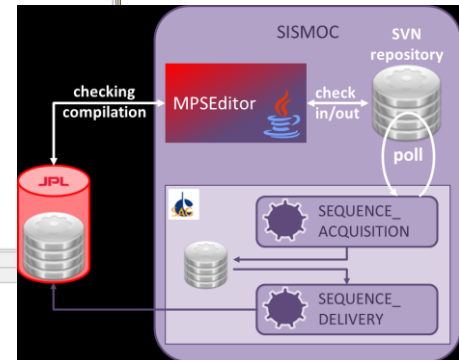


Fig. 22: MPSEditor integration into SISMOC

MPS Editor can actually be used for any kind of sequence for they are all written in the same programming language. During the development phase of the mission, it was actually the only way to generate event sequences. Yet those sequences are very error-prone because of the numerous touchy details that the operator must take care of, such as the date conversions to LOBT or AOBT. Fortunately, the SISMOC embeds a dedicated tool presented section 4.4.1, which automates the generation of these sensitive sequences.

4.4 Event Requests

Section 3.1 mentioned the impacts of the limited bandwidth of the Mars-to-Earth communication link on the operational organization and the ground system. The FSW FIR filters are pre-configured on-board, based on scientific specifications, to produce a list of filtered and decimated data (see sections 3.2 and 4.2). The resolution of these data was chosen as a trade-off to match two requirements:

- make it possible for the science team to identify interesting seismic events,
- cover the full period of data acquisition by the instruments despite the limited downlink transmissions.

Because of the latter requirement, these data constitute the so-called “continuous data” flow (see section 7.2 of [5]). After these continuous data are processed and published by the SISMOC, the science team analyze them thoroughly. When they identify some interesting seismic events, they may wish to get higher resolution data on the event time window to better characterize its nature, origin, strength, etc. These requested higher resolution data constitute the so-called “event data” flow (see both flows on Fig. 24). A part of the Mars-to-Earth communication bandwidth is allocated to this flow in such a way that




Fig. 24: Example of continuous and event data

there is always room enough for the continuous data flow too. Unfortunately, it was predictable that the science team would soon ask for more data than allowed by the limitation on the event data bandwidth. It was also predictable that sometimes there would be issues with orbiters that would result in less bandwidth than expected. It was thus decided that the event requests would be classified according to three levels of priority: low, medium and high. Whenever an event sequence runs on-board, the SEIS FSW extracts the event data from the FRB, applies its FIR filters in a different configuration than for the continuous data (obviously, since the decimation ratio is not the same !) and sends them to the lander FSW with their priority attached. The FSW builds the space packets of event data and stores them into three send buffers, one for every priority level. Every time the communication is established between the lander and an orbiter, these buffers are emptied oldest data first, starting with the high-priority buffer – see *Fig. 25*.

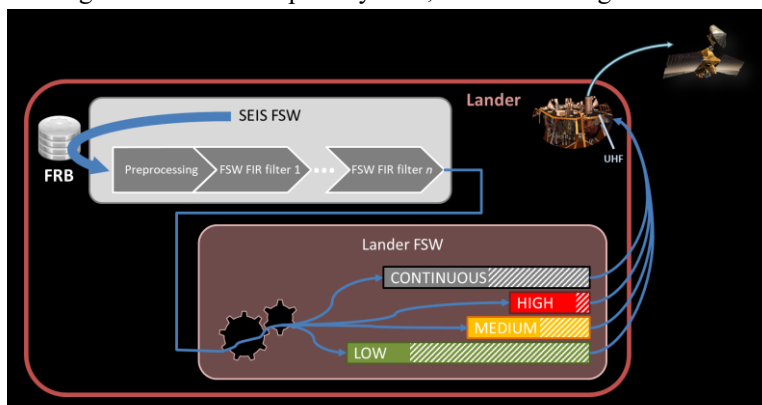


Fig. 25: Event buffers on-board

The event requests that are uploaded on the lander for execution are the result of a complex selection process that relies on various estimations:

- 1) the actual available communication bandwidth when the event request is expected to run on-board,
- 2) the fill level of the on-board event buffers at the same time: it depends on how many data were requested before and how many could be actually sent to an orbiter during the previous passes,
- 3) the availability of the requested data in the FRB, for they may have been overwritten by the time when the sequence runs,
- 4) the amount of data produced by an event request: since some data compression is involved here, an exact prediction is impossible,
- 5) the duration of the event request on-board processing: it has to execute within a limited time range fixed by the activation of the lander, right before an orbiter pass.

Points 1) and 2) are under the responsibility of JPL. Points 3) to 5) are on CNES, as well as the generation of the event sequences. The prioritization of the event requests is on the science team, but they need help in making those decisions accordingly to all the constraints and without wasting the available bandwidth. For this purpose, CNES has developed in the SISMOC scope a set of two web applications, EBM and ERP Tool, which run on a server open to the internet and both provide a web interface protected by a personal account authentication mechanism.

4.4.1 Event Buffer Management and Event Request Proposal Tool

EBM relies on a model of the status of the FRB and the event buffers. Its purpose is the checking for the availability of the requested data before an event request is planned. Once this is validated, EBM can also generate the event sequences. Its web interface is mostly destined to the instrument team for the follow-up of the buffer statuses. Yet, a few months after the landing, this interface was almost only used for its sequence generation feature.

The ERP Tool is also composed of a server and a web client interface. It is the entry point when it comes to the event request life. It is elaborated enough to take into account all the points and limitations above. The counterpart is that only people educated to the operational process and the mission specificities are able to use it correctly. The access to it is moreover limited to a close circle of scientists involved on the mission. These are the reasons why the direct CNES scientific partners on this mission, IPGP and ETHZ, have developed their own interface for the submission of ERP, based on the ERP Tool REST API [10]. IPGP has even developed a complete portal for the submission and the follow-up: the SEIS Data Portal (SDP). Any scientist can ask for an account on this portal and start submitting ERP.

4.4.2 ERP submission

A submitted ERP contains mainly data channels (among the output channels of the FSW FIR filters) and the time interval of the data to be requested. The user can also set a priority which defaults to medium. When it receives an ERP submission, the ERP Tool first checks if

- 1) the requested data are still in the FRB (for it is a ring buffer as mentioned section 3.2),
- 2) the data channels were produced within the requested time interval.

Check 1) requires the ERP start time be compared to the date of the oldest data in the FRB. That date is reported by some specific TM called ancillary data. The SISMOC acquisition process for ancillary data forwards them to the ERP Tool which extracts from them that FRB date and stores it in its local database. Yet, since the instruments produce data continuously, this may not render the exact FRB status at the very time when the TM are processed on-ground, nor at the time when the ERP is submitted, and it is even more obsolete at the time when the ERP is expected to be processed on-board. Predicting the exact FRB status at a date T would require a perfectly accurate knowledge of the evolutions of the data production rate from the time when the TM was emitted until T . This production rate depends on the instrument configuration that operational actions may alter. If the prediction is theoretically possible, it is practically unachievable because it is very difficult to foresee accurately the moments when the configuration changes. Yet the accuracy of the prediction is not essential here because it is always possible to keep some margins in the ERP requested dates and because the impact of requesting data that are no longer in the FRB is of less importance: an error report and a waste of bandwidth are no mission critical issues. Besides, the repetitive scheme of the operations makes it possible to predict the FRB status with a suitable precision based on its past evolution. Therefore, the ERP Tool implements some simple heuristics for the estimation of the FRB oldest data date. When the scheme of the operations changes however (for instance when it was decided to acquire VBB data at a higher rate or when the instruments were switched off regularly to spare energy), the model must be adapted and this involves all teams again. The science team are always involved in the decision of changing the scheme of operations, even when it is due to energy constraints. The instrument and GDS teams work then together to establish the new ERP Tool prediction model of the FRB status and correct it after some weeks if need be.

Check 2) is much more reliable. The production of data channels is bound to the sensor configuration (power status, gain, acquisition rate, etc.). Only a subset of the defined channels is thus available for downlink at a given time period. Section 3.3.1 presents how the sensor configuration changes are translated into lists of active miniSEED channels, on the way to produce the dataless file. In a similar way, the SISMOC process *EBM_TM_STATUS* generates text files that reflect all changes in the status of the sensors. Thanks to these files, the ERP Tool stores in its local database a representation of the history of sensor statuses. This makes it straightforward to check if the channels of a submitted ERP were produced within the requested dates.

Unfortunately, there are some cases when an instrument is powered OFF and the TM does not render that status. When APSS and SEIS switch to SAFE mode, neither the PAE nor the E-Box have time to emit the HK packet that would mark the sensors as OFF. The PAE has also a design issue when it is switched OFF: all sensors are shutdown sequentially but the status OFF of the last one is not reported in the HK TM because the PAE is OFF before sending it. Because of this, the ERP Tool may mistake the sensors as ON when they are not and wrongly accept requests on the OFF period. An operational workaround has then been settled. Just as the instrument team know of the power ON dates of the instrument (see section 3.3.3), they know of the power

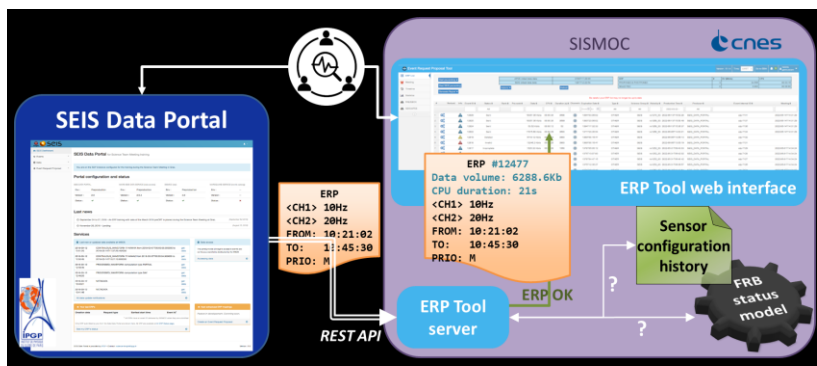


Fig. 26: ERP validation at submission time

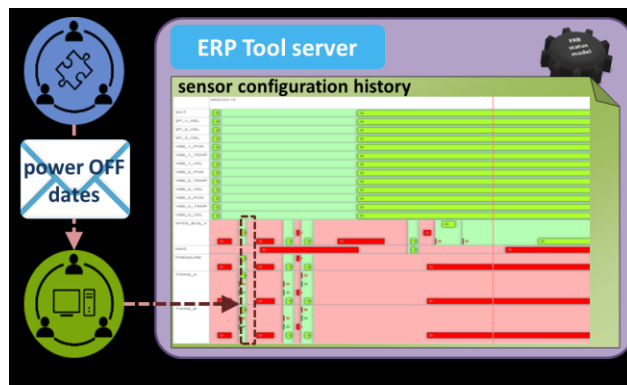


Fig. 27: ERP Tool - Manual update of sensor history

OFF dates, either planned or unexpected, when they analyze the TM. They inform then the GDS team who manually modify the sensor configuration history into the ERP Tool database (see Fig. 27).

Once a submitted ERP is validated against checks 1) and 2), the ERP Tool computes an estimation of the request processing duration and of the data volume that it would produce for the download. The ERP is then recorded in the local database with a unique ID and the status PROPOSED.

4.4.3 ERP selection

The operational organization includes a set of weekly meetings for the ranking of all submitted ERP, which ends with the Event Selection Meeting (ESM). This last meeting gathers a science leader, the Science Operations Working Group chair, the CNES instrument and GDS teams. It aims at defining the final set of ERP that can be planned for an execution on-board, *i.e.* which can be converted into actual event requests. In order to propose a list of ERP that could match all the constraints, the ERP Tool require the following (see Fig. 28):

- 1) the various time windows allocated to the execution of event request sequences – these time windows are called “wake-ups” because they take place when the lander wakes up to process communications and/or run command sequences,
- 2) the free space in the various event buffers at the beginning of every wake-up (this value is already moderated by the JPL GDS team with regards to the estimated available bandwidth),
- 3) the estimated oldest data available in the FRB at the beginning of every wake-up,
- 4) the dates of the requested data, the rank, the estimated processing time, the estimated data volume and the priority level of every ERP.

Points 1) and 2) are provided by the JPL GDS team in a Science Activity Plan that is the central product of the uplink operational phase. It comes in several versions as it is constantly modified during that phase until it is ready for the upload. The CNES instrument team are responsible for providing the right version to the ERP Tool prior to the ESM. Point 3) has been thoroughly discussed in section 4.4.2. It has already been said that a wrong estimation was not critical for the mission. Therefore, if the ERP Tool detects a risk that some ERP request data that may not be available at execution time, then it simply raises a warning on that ERP to the attention of the ESM actors.

Point 4) is already managed by the ERP Tool, computed at submission time and stored in its local database.

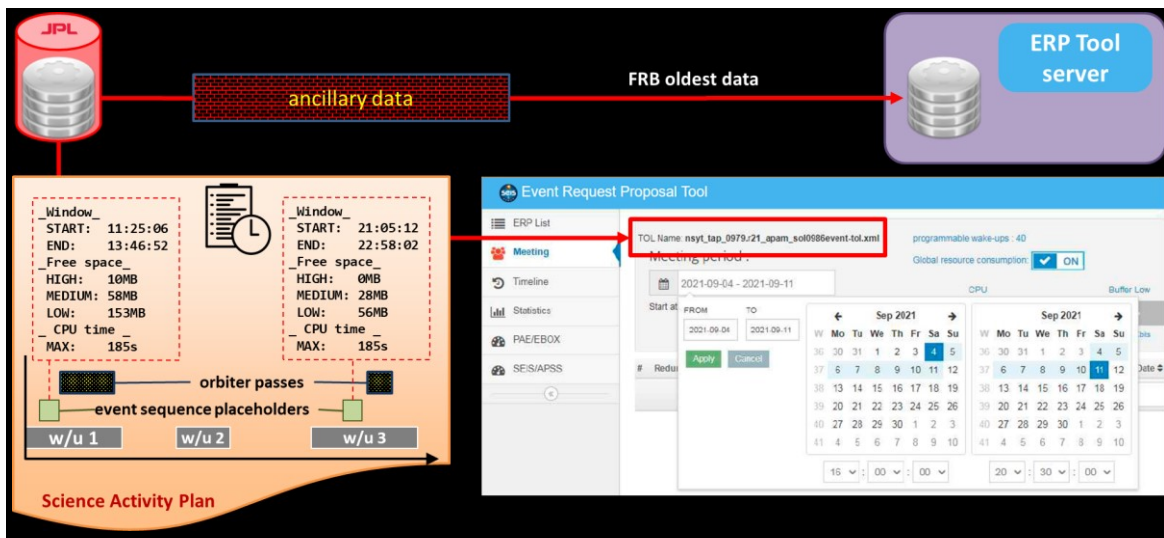


Fig. 28: Event Selection Meeting preparation

When the ESM manager opens the meeting, the ERP Tool performs a first ventilation of the ERP. Every ERP starting from rank 1 is ventilated in the first wake-up where there is enough room accordingly to the free space in the event buffer that matches the ERP priority. If needed, an ERP may be split across several wake-ups. Experience has revealed that the processing duration was never an issue, even when the energy level was so low that the wake-ups had to be shortened. Experience has also shown that this first ventilation was inefficient most of the time, because of the priority level constraint.

Then the ESM manager uses a specific feature of the tool: the automatic buffer allocation (see Fig. 29). The underlying principle of this ventilation algorithm is that having a high priority event request sent to the medium event buffer is always better than planning a medium event request in its stead, just because there was not enough room in

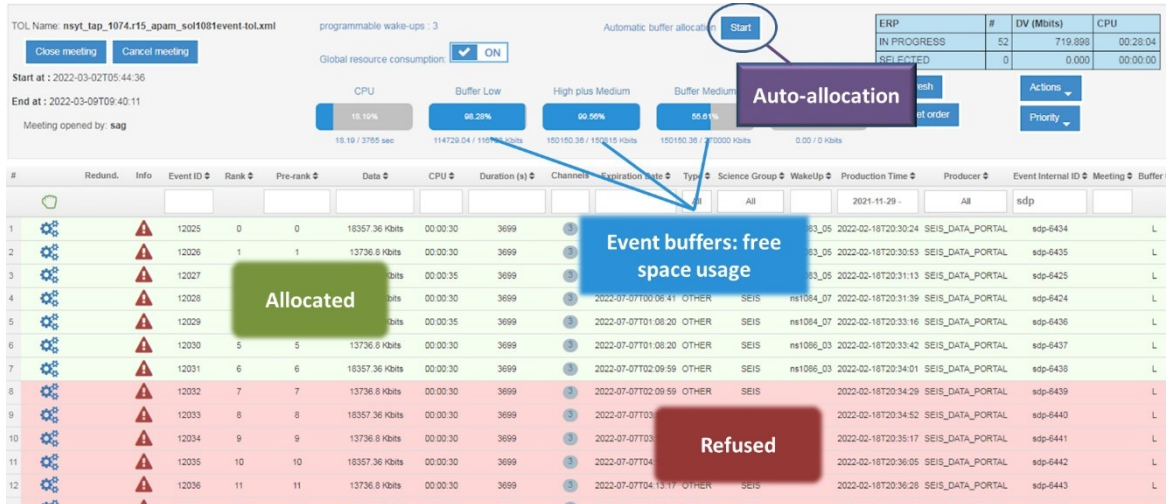


Fig. 29: ERP Tool - Event Selection Meeting view

the high event buffer. Obviously, the same goes for medium event requests in the event low buffer. Because the ERP were most of the time submitted with the default medium priority and because the event low buffer is much bigger than the other two, this second ventilation has shown outstanding results since the beginning of the operations. There have been a few manual attempts to improve the usage of the buffers (and through that, the usage of the available bandwidth): people would change ranks or priorities on some ERP during the meeting. They never ended in a better result than the automatic buffer allocation, so that people just stopped trying and started submitting ERP with smaller amounts of requested data.

Once all participants agree on the selected ERP list, the manager closes the meeting. The allocated ERP switch to the status PLANNED whereas the unallocated switch to the status POSTPONED, because they will still be candidates at the next meeting. The instrument team import then the meeting results into EBM (actually, both tools share the same database), check every wake-up for potential issues on allocations and eventually generate the event request sequences. The SAG (see section 2) retrieves automatically these event sequences to let the operator launch the built-in delivery-to-JPL process.

Once they are delivered, the sequences are also sent to the EBM server. The FSW command that retrieves the event data requires an ERP ID as an argument, the very same unique ID that was set at submission time. It is thus straightforward to identify from the sequences which ERP are to be switched from the status PLANNED to SENT.

4.4.4 ERP follow-up

As it was mentioned above, the ERP ID is used on-board when retrieving the event data. It is then abusively called the “event ID”. It follows the data in the TM packets down to Earth and back to the SISMOC. Hence the SISMOC, through its process *EBM_TM_STATUS*, informs the ERP Tool of the amount of data downloaded for any sent ERP. The ERP Tool keeps track of the total volume of data previously received for every ERP. As soon as this volume is greater than 0, the ERP switches to the status DOWNLINK_IN_PROGRESS, and when it equals (or crosses) the initially estimated data volume, then the ERP switches to the status COMPLETE. Naturally this status may never be reached. Loss of TM packets, a wrong initial estimation of the data volume or even data no longer in the FRB would all lead to this situation. If an ERP remains too long in the status DOWNLINK_IN_PROGRESS then it switches to the status INCOMPLETE.

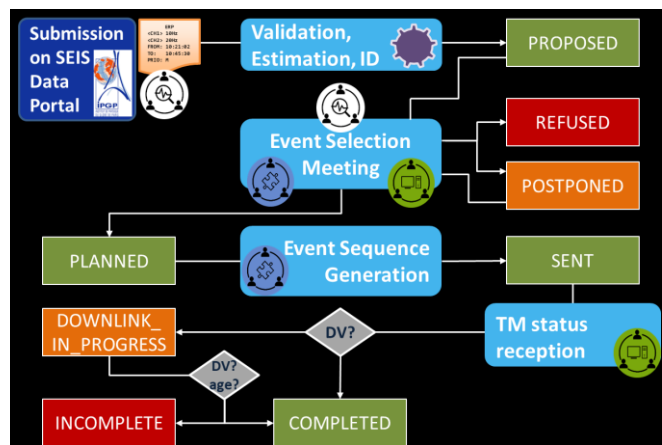


Fig. 30: ERP state machine

5 Conclusion

The achievements of the InSight mission and its SEIS seismometer have been highly and widely appreciated. Obviously, this is primarily due to the outstanding quality of the instrument. Put aside the amazingly successful management of the cruise, the landing and the instrument deployment by the JPL, such quality would have been pointless if the operational context, from the lander down to the GDS, had not been able to match that quality level in the final scientific products. This could only be achieved through a close collaboration between the science, instrument and GDS teams. This paper intended to emphasize how much the SISMOC, developed at CNES as the SEIS and APSS ground system, was a central point of this collaboration, both as an environment shared by all the teams and as the destination of some configuration workflows.

The monitoring of the instrument with the generation of state-of-health reports, the selection and follow-up of the event data requests and the generation of commanding sequences are all ground activities obviously closely linked to the SISMOC. It was less obvious however that the configuration of the FIR filters sets delays to be taken into account in the dating of the data, not obvious that some on-board commands may be only destined for the on-ground processing or that the instrument power cycles must be reported on the ground time-correlation tool configuration as well as on the event request management tool. All these activities are generally considered as tasks for the operational engineers and yet they have immediate consequences on the ground system. They must thus be driven in collaboration with the GDS team, since the early stages of the mission development.

Acknowledgements

Special thanks to R. Llorca-Cejudo, CNES, Toulouse, whose slideshows helped a lot in understanding the on-board specificities of the lander and the instruments.

Many thanks to all the people involved in the design and development both of the instrument and the SISMOC, to JPL, CNES, IPGP and all the InSight partners for this wonderful mission.

Last but not least, reserved acknowledgements to C. Agard and T. Bloch, CNES, Toulouse, for their commitment in maintaining the SISMOC in good operational conditions.

References

- [1] NASA, "Discovery Program," [Online]. Available: <https://www.nasa.gov/planetarymissions/discovery.html>.
- [2] C. Krause and al., "HP3 – Experiment on InSight Mission – Operations on Mars," in *16th International Conference on Space Operations*, Cape Town, South Africa, 2021.
- [3] CNES, "InSight mission releases data on monster marsquake," CNES, 20 10 2022. [Online]. Available: <https://presse.cnes.fr/en/insight-mission-releases-data-monster-marsquake>.
- [4] T. Kawamura and et. al, "S1222a-the largest Marsquake detected by InSight," *Geophysical Research Letters*, 2022.
- [5] P. Lognonné, W. B. Banerdt and D. Giardini, "SEIS: Insight's Seismic Experiment for Internal Structure of Mars," *Space Science Reviews*, vol. 215, no. 12, 2019.
- [6] ESA, "ExoMars/TGO," [Online]. Available: https://www.esa.int/Enabling_Support/Operations/ExoMars_TGO_operations.
- [7] T. K. Ahern and B. Dost, "SEED Reference Manual," March 2006. [Online]. Available: http://www.fdsn.org/pdf/SEEDManual_V2.4.pdf.
- [8] ETHZ, "QuakeML," [Online]. Available: <https://quake.ethz.ch/quakeml/>.
- [9] C. M. Pilato, B. Collins-Sussman and B. W. Fitzpatrick, *Version Control with Subversion*, 2nd edition, O'Reilly, 2008.
- [10] R. T. Fielding, *Architectural Styles and the Design of Network-based Software Architectures*, vol. Chapter 5: Representational State Transfer (REST), Irvine: University of California, 2000.