Science Data Volume management for the Rosetta spacecraft

M. Pérez-Ayúcar[1], M. Almeida[2], M. Ashman[2], M. Costa[2], J. García[2], R.M.T. Hoofs[3], M. Kueppers[4], D. Merritt[2], J. Marin[2], F. Nespoli[2], E. Sanchez[2]
European Space Astronomy Center (ESAC-ESA), Villanueva de la Cañada, Madrid, Spain

and

S. Chien[5], G. Rabideau[6]
Jet Propusion Laboratory (JPL-NASA), Pasadena, USA

The International Rosetta mission is a European Space Agency (ESA) mission with contributions of its member states and National Aeronautics and Space Administration (NASA). Launched on 2nd March 2004, and headed to rendezvous the comet 67-P/Churyumov-Gerasimenko, it comprises the Rosetta orbiter and the Philae lander. The main scientific goals are to map the comet 67-P/Churyumov-Gerasimenko by remote sensing, to in-situ examine its environment and to study its evolution in the inner solar system. The spacecraft arrived at the comet on 6th August 2014 and Philae was deployed on the comet´s surface on 12th November 2014. It is the first spacecraft to ever rendezvous with a comet and orbit it. It also placed for the first time a lander onto the surface of a comet. On 13th August 2015 Rosetta passed perihelion. The mission is expected to continue its science until fall 2016 when Rosetta might be placed onto the surface of the comet. The Rosetta science operations are coordinated by the Rosetta Science Ground Segment, located at ESA´s European Space Astronomy Centre (ESAC) in Madrid, Spain. The Science Ground Segment has the role to construct, in coordination with the Science Teams and the Mission Operation Centre, a plan of scientific observations and provide conflict free operational timelines for upload and onboard execution. The data volume allocation and management is a key part on the final design of the Rosetta operations. In this paper we will describe how the available scientific data dump capacity is managed for a successful science data return to Earth. The overall objective is to fully utilize the ground stations capability, balancing the data dumps for a low latency, while avoiding any onboard data losses by memory overwriting. We will focus mostly on the uplink processes, from the initial data capacity estimation and initial allocations to the teams, to the final iterations and modifications for command upload. We will stress some particularities and challenges imposed by the very nature of the unpredictable environment of the comet. We will finally describe briefly some of the analysis done during actual observations execution and dump to ground. This info can then be fed back in a short-term loop into the next planned science observations.

Nomenclature

AU = Astronomical Unit
BoT = Beginning of Transmission
DSA = Deep Space Antenna
DSN = Deep Space Network
DV = Data Volume
EDF = Experiment Description File
EoT = End of Transmission

---

[1] Rosetta Operations and Planning Group Leader, Science Operations Department, ESA/ESAC, Madrid, Spain
[2] Rosetta Operations and Planning Group Engineer, Science Operations Department, ESA/ESAC, Madrid, Spain
[3] Rosetta Operations and Planning Group Advisor, Science Operations Department, ESA/ESAC, Madrid, Spain
[4] Rosetta Science Ground Segment Uplink Coordinator, Science Operations Department, ESA/ESAC, Madrid, Spain
[5] ASPEN/Rosetta Lead, JPL/NASA, Pasadena, USA
[6] ASPEN/Rosetta Software Lead and DALLOC Lead, JPL/NASA, Pasadena, USA

EPS = Experiment Planning System
ESA = European Space Agency
ESAC = European Space Astronomy Center
ESTRACK = European Space Tracking Network
FIFO = First In First Out
FCT = Flight Contorl Team
FECS = Flight Events and Communications Skeleton (file)
FD = Flight Dynamics
GS = Ground Station
HGA = High Gain Antenna
HK = House Keeping
HSL = High Speed Link
MTP = Medium Term Planning
MAPPS = Mission Analysis and Payload Planning System
NASA = National Aeronautics and Space Administration
LTP = Long Term Planning
LS = Liaison Scientist
OBDH = Onboard Data Handling Bus
ORCD = Operations Requirements and Constraints Document
OWLT = One Way Light Time
PI = Principal Investigator
POR = Payload Operations Request (file)
PDOR = Payload Direct Operations Request (file)
PTR = Pointing Request (file)
RMOC = Rosetta Mission Operations Center
RSGS = Rosetta Science Ground Segment
SSMM = Solid State Mass Memory
STP = Short Term Planning
TLM = Telemetry
VSTP = Very Short Term Planning

I. Introduction

The Rosetta Science Ground Segment (RSGS) primary role in the Rosetta project is to support the Project Scientist and the Science Working Team, to ensure the coordination, development, validation and delivery of the desired science operations plans and their associated operational products. RSGS also provides support to the Principle Investigator (PI) teams in order to ensure the provision of adequate data to the Planetary Science Archive. A summary of the eleven science instrument packages on the Rosetta Orbiter, and the ten instruments on the Philae Lander, is shown in the Appendix.

Within the RSGS, the Operations and Planning Group[1] is responsible for the Medium, Short and Very Short Term cycles (MTP, STP and VSTP), a stepped sequential process that implements the higher level science Long Term Plan (LTP), developed by the RSGS Liaison Scientist Group[2], into the final science Pointing Requests (PTR) and Payload Operational Requests (POR), the command sequences to be executed on-board the spacecraft.

The Operations and Planning Group takes the high level timeline of science observations and, via the established tools, interfaces and planning processes, works

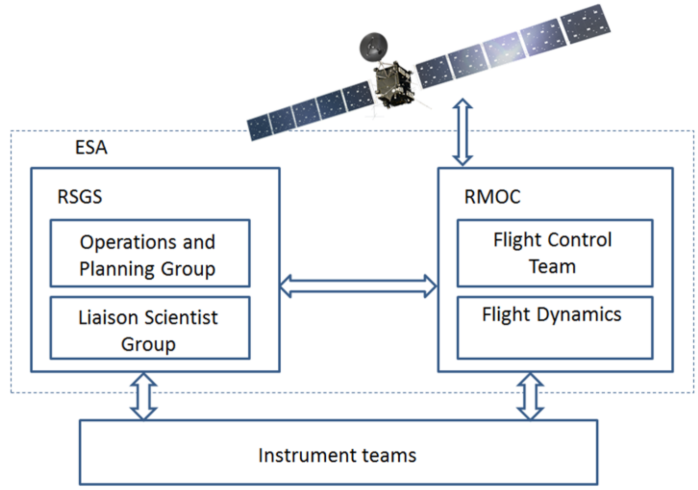

Figure 1. Rosetta planning interfaces. *The Rosetta main planning teams are the RMOC (Rosetta Mission Control Center), the RSGS (Rosetta Science Ground Segment) and the Instrument Teams.*

with the Orbiter PI teams (and also the Lander Ground Segment) to generate and validate the associated operational products. These are ultimately delivered to the Rosetta Mission Operations Center (RMOC) for final verification before upload to the spacecraft for execution.

Rosetta science planning occurs in a highly constraint environment with limited resources. It has to comply with the severe constraints imposed by its unique design characteristics, the spacecraft and instrument capabilities, and the extreme cometary environment. The available resources to achieve the Rosetta goals are also limited and shared among the instruments under competition.

Some of the most important Rosetta constraints are the pointing capability, the Sun illumination avoidance on certain sensitive panels, the instrument constraints (such as Sun avoidance on field of views, the thrusters actuation avoidance, and the cross interference), and the power and data downlink available. The latter two are dependent on the large distance variation due to the large eccentricity of the comet orbit. This paper will focus on how the data volume resource is managed to achieve an optimized data downlink that meets the scientific goals of the mission.

II. General description of Rosetta data downlink aspects

A. Ground station resources

Rosetta relies primarily on the ESA ESTRACK Deep Space network (34m antennas in Cebreros, New Norcia and Malargüe). NASA provides extra coverage support to fulfill Rosetta's data needs with the Deep Space Network (DSN, 35m and 70m antennas in Madrid, Canberra and Goldstone). The Flight Control Team negotiates with other missions which ground stations are allocated for Rosetta, based on the scientific capacity needed. Rosetta books different ground stations around the globe to ensure an almost continuos coverage per day. The dump capacity (stations used, pass durations and bit rate) is communicated to the Science Ground Segment via the Flight Events and Communications Skeleton file (FECS). The Rosetta project regularly requests NASA 70m antennas support to boost the downlink performance and ensure high rates during critical phases of the mission. ESTRACK allocations are confirmed with ~6month lead time, whereas DSN allocations are only confirmed ~2 months before execution.

Figure 2. Rosetta ground station facilities around the globe. *Rosetta dumps the mission data mainly through the ESA-ESTRACK Deep Space Antennas (DSA, in blue), with important support from NASA Deep Space Network (DSN, in red). Their distributed location around the globe ensure a ~24h coverage.*

The telemetry downlink bitrate is a function of the distance to Earth and the size of the antenna used at the groundstation, as defined in the link budget. The highly elliptical orbit of the comet Churyumov-Gerasimenko (heliocentric distances from 1.2 to 5.6 AU, geocentric distances from 1.5 to 5.5 AU) means that the data volume available varies substantially over the Rosetta mission phases as illustrated in Figure 3.

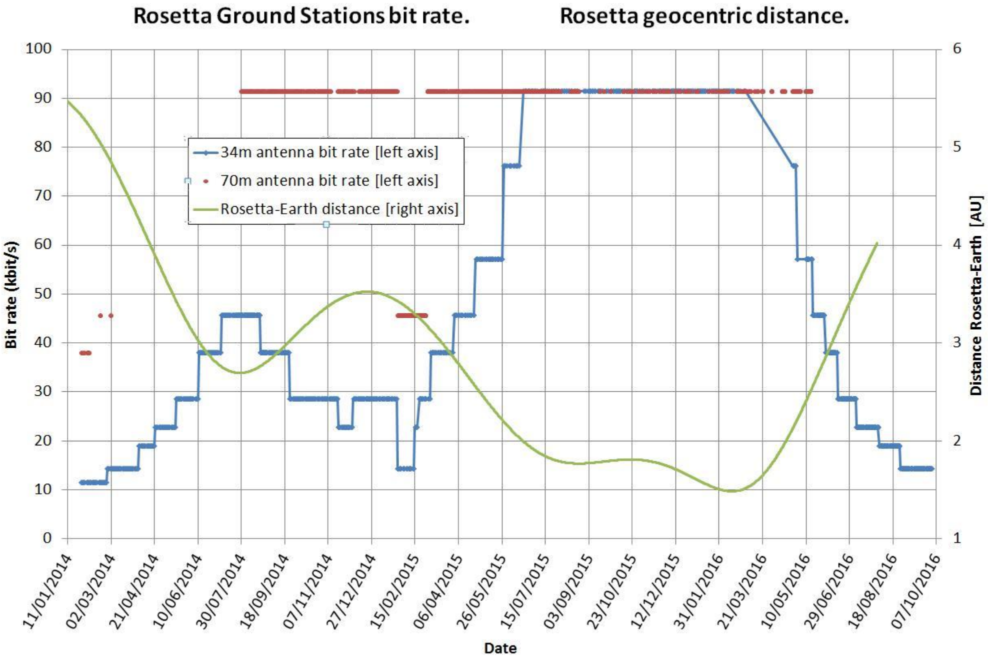

Figure 3. Rosetta ground stations bit rate and Earth distance. *The left axis compares the actual bit rate (kbit/s) used during the Rosetta mission for the ESTRACK/DSN 34m antennas (in blue), and the DSN 70m antennas (in red dots). The right axis shows for reference the Rosetta-Earth distance in astronomical units (AU). The lowest bitrates, apart from the early 2014 hibernation exit, happened around February 2015 due to a solar conjunction. The maximum onboard telemetry rate capacity is 91.412 kbps.*

B. Payload and Platform data

Rosetta handles two main types of telemetries for downlink to ground.

1) The Payload data comprises all the science and housekeeping (HK) data gathered from the instruments.
    - The Science data is the science telemetry during the observations
    - The instrument housekeeping data is a small amount of auxiliary information that is generated to understand the status of the instrument and the conditions of the science measurements. The housekeeping data is modelled with a worst case value taken for the entire duration of the instrument being on, as shown in Table 1. The total amount modelled is a constant 1500bps.
2) The Platform telemetry are the data gathered by the spacecraft subsystems and the navigation cameras
    - The house-keeping data comprise the Data Management Subsystem, the Attitude and Orbit Control Management Subsystem, and the memory cell acquisition. The overall rate is around 16.2MBytes/day, or a constant bit rate of around 1500bps. This amount can vary according to the needs of the mission phase.
    - The Navigation images are acquired usually during navigations slots. For distances to the comet typically 100km or lower, 4 images are taken in a 2x2 raster. For distances higher than 100km, a single image in a nucleus stare pointing is taken. There is usually one navigation slot every 5h, and each image is 1.61 Mbytes.

C. Spacecraft memory capacity

All the Rosetta payload is stored in a physical memory unit called the SSMM (Solid State Mass Memory). The SSMM is partitioned onboard to allocate independent data areas for each of the instruments and for the platform, according to their needs. The Rosetta science partitions are determined in the Table 2.

Instrument	HK SCI (ORCD)	
ALICE	121.6	bps
CONSER	150.0	bps
COSIMA	26.0	bps
GIADA	100.0	bps
MIDAS	147.2	bps
MIRO	103.0	bps
ROSINA	70.0	bps
RPC	54.0	bps
SR_SCI [OSIRIS]	19.0	bps
VIRTIS	250.0	bps
LANDER	108.0	bps
TOTAL	**1148.80**	bps

Table 1. Rosetta Science Housekeeping: maximum production. *The numbers provided represent a worst case production per instrument, in bits per second, as in the ORCD (Operations Requirements and Constraints Document)[3].*

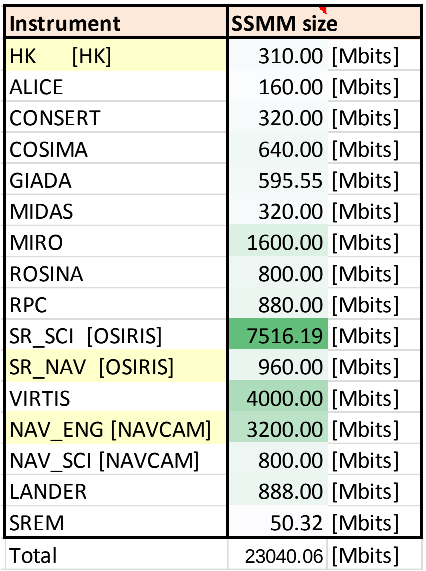

Instrument	SSMM size	
HK [HK]	310.00	[Mbits]
ALICE	160.00	[Mbits]
CONSERT	320.00	[Mbits]
COSIMA	640.00	[Mbits]
GIADA	595.55	[Mbits]
MIDAS	320.00	[Mbits]
MIRO	1600.00	[Mbits]
ROSINA	800.00	[Mbits]
RPC	880.00	[Mbits]
SR_SCI [OSIRIS]	7516.19	[Mbits]
SR_NAV [OSIRIS]	960.00	[Mbits]
VIRTIS	4000.00	[Mbits]
NAV_ENG [NAVCAM]	3200.00	[Mbits]
NAV_SCI [NAVCAM]	800.00	[Mbits]
LANDER	888.00	[Mbits]
SREM	50.32	[Mbits]
Total	23040.06	[Mbits]

Table 2. SSMM partitions. *The numbers provided state the allocated space in the SSMM for each instrument packet store. The memory sizes are coded with an increasing dark green background. Platform data stores have a yellow background.*

D. Instrument internal memories

The Osiris instrument is the only one in Rosetta that uses an internal memory to temporarily store the acquired science data. Osiris can store the data in several internal queues, and then command the dump to the spacecraft for subsequent downlink to Earth. This feauture provides higher flexibility to buffer dumps in periods of very high acquisition (like lander search campaigns), for multi-instrumnet simultaneous observations (such as close flybys), during low dump capacity periods (as in solar conjunction), or to selectively choose which data to dump.

E. Data link from instrument to the SSMM

The instruments use the on board data handling bus (OBDH) to transfer the science and housekeeping data to the spacecraft. OSIRIS, VIRTIS and the Navigation camera have a dedicated interface bus, called the High Speed link (HSL), as their acquisition data rates are higher than the rest of the experiments. The traffic over the OBDH should be less than 84000 bps, whereas on the HSL can be up to 10 Mbps.

F. Prioritized dump

Rosetta has the capability to select priorities for each data store during each ground station pass. The stored spacecraft Housekeeping (HK) telemetry and the camera images used for navigation are always dumped first, as they are critical for mission safety. The Rosetta Science Ground Segment implements the priority dump management of the science data. The prioritized dump schema is sent to the Flight Control Team to upload to the spacecraft.

III. The data volume management cycle

The Rosetta Science Ground Segment is responsible for the management of the Rosetta downlink, with a particular attention to the science data generation, memory stores control and priority dumping. The RSGS ensures that all the science data requested by the experiment teams can be generated and downlinked with the available resources around the time of the execution of the scientific observations. The iteration cycle in which the data volume is managed will be explained in this chapter. It will cover sequentially the initial downlink capacity estimation, the command iterations with the PI teams, the last minute modifications close to execution, and the post-execution analysis.

A. Available data volume estimation from the FECS file

The link budget is calculated by the Mission Control Team, as a function of the distance to Earth and the size of the antenna used at the groundstation. After negotiation with other missions, the pass duration and the telemetry bit rate is communicated to the Science Ground Segment via the Flight Events and Communications Skeleton file (FECS). A new FECS file is updated and re-generated every week. It is automatically checked by the Operations and Planning Group at reception with several analysis tools that summarize the future station coverage, plus the eventual modifications to the already existing bookings. An example of the output of one of the analysis tools is shown in Table 3.

MTP	ESA Pass Duration (h/day)	ESA Dump Duration (h/day)	DSN 34m Pass Duration (h/day)	DSN 34m Dump Duration (h/day)	DSN 70m Pass Duration (h/day)	DSN 70m Dump Duration (h/day)	Bistatic Radar/Limb Sounding passes (h/day)	Total Data Dump (bits)
MTP_001	9.7689	8.3028	4.4762	3.1441	0.0000	0.0000	0.0000	1.30566E+10
MTP_002	8.7400	7.4353	5.4531	4.1500	0.0000	0.0000	0.0000	2.67279E+10
MTP_003	10.7297	9.5485	5.4923	4.3135	0.0000	0.0000	0.0000	3.76323E+10
MTP_004	12.0964	10.9972	4.9511	3.8524	0.0000	0.0000	0.0000	5.34955E+10
MTP_005	12.7516	11.3159	4.3137	3.3809	0.4075	0.3047	0.0000	7.37687E+10
MTP_006	16.5138	13.5997	1.3750	1.0490	4.6875	3.7058	0.0000	1.10922E+11
MTP_007	17.3167	15.1267	0.1627	0.1109	4.6071	3.6706	0.0000	6.92432E+10
MTP_008	17.8946	15.6558	0.3575	0.2457	3.6022	2.7064	0.6452	7.85261E+10
MTP_009	18.6301	16.2395	0.4202	0.2925	3.3527	2.5014	0.0000	7.14398E+10
MTP_010	15.1513	13.0485	0.7917	0.5223	2.9137	2.1980	0.7321	5.80675E+10
…	…	…	…	…	…	…	…	…
MTP_028	10.2716	8.4678	6.5536	5.2591	0.7798	0.5951	0.0000	1.23058E+11
MTP_029	14.5297	11.8761	0.0000	0.0000	0.0000	0.0000	0.0000	6.74753E+10

Table 3. Ground station passes analysis. *Pass and dump durations from a weekly FECS file. The analysis starts at the beginning of the Rosetta comet phase (MTP001), and is split in MTP periods (column 1). Each MTP last 4 weeks (28 days). For reference, MTP001 started on 17th March 2014. MTP029 started on 3rd May 2016. MTP011 to 27 are skipped for conciseness. The duration statistics are provided for ESA 34m ESTRACK stations (columns 2 and 3), NASA DSN 34m support (columns 4 and 5) and NASA DSN 70m support (columns 6 and 7). Special campaigns for bistatic radar and limb sounding are in column 8. Total data dump (in bits) is computed in column 9, using the available bit rate as in Figure 3.*

Typical Rosetta pass durations are ~18h per day, with effective dump durations of ~16h per day. The total amount of data dump per day ranges from 2 to 6 Gbit per day.

Please note that the full pass duration is slightly longer than the actual data dump duration. This is due to the time needed to establish the link and some added margins. For example, the onboard dump start time happens within 12 minutes from the first radiated signal at the ground station Beginning of Track time (BoT), plus the one-way light time (OWLT). These 12 minutes take into account the sweeping and lock of the onboard communications carrier. The end of the dump is terminated onboard one light travel time and 5 min (margin) before the ground End of Track time (EoT).

B. Initial data volume allocation per instrument

At the beginning of each medium term planning period, the Liaison Scientist group produces a high level science harmonization to allocate the available observation slots (see [2]). The data volume is also assigned, based on the ground station availability at the time. The ground station bookings extends ~6 months in the future for ESTRACK ground stations, and ~2 months for DSN. The booked period is usually long enough to cover the cycle in progress. In this case, the Operations and Planning Group provides to the Liaison Scientist Group the expected science data volume available, after subtracting from the FECS estimations the navigation images, the housekeeping telemetries and the high gain outage periods during orbit control propulsive maneuvers (OCM). Typically there are two ~6h duration OCM slots every week. Note that if the data volume is not yet available at the start of the planning period, it can estimated from the expected ground station bitrate, with an assumption of the effective downlink time (around 16 hours per day).

The effective Science data volume can reach up to 600 MBytes per day, during high rate periods, and 150Mbytes per day with low bitrates. In case of reduced data volume periods, like conjunctions, DSN 70m antenna support is requested from DSN.

The data volume is allocated per instrument, according to their level of resource usage. Some instruments generate such a small quantity that any requested amount is granted (CONSERT, COSIMA, GIADA, MIDAS, PHILAE, and SREM). Other instruments (ALICE, MIRO, ROSINA, and RPC) typically generate a significant near continuous amount of data, and their requests are accepted as envelopes restricted during some mission phases. Last, OSIRIS and VIRTIS produce large amounts of data in bursts (image acquisitions), and they are given ~70% of the available volume.

An example of the instrument allocations for a typical planning period (MTP020) is shown in Table 4.

Planning period	MTP015 week 1-2 (08-21 April 2015)	MTP015 week 3-4 (21 April-05 May 2015)	MTP015 Total
Science DV available [Mb]	63557	68006	**131563**
Instrument	DV Allocation [Mb]	DV Allocation [Mb]	**DV Allocation [Mb]**
OSIRIS	33776	30555	**64331**
VIRTIS	19764	17248	**37012**
MIRO	2814	2812	**5626**
ALICE	2990	3042	**6032**
ROSINA	2272	2271	**4543**
RPC	6053	6050	**12103**
GIADA	87	87	**175**
MIDAS	388	388	**776**
COSIMA	594	594	**1188**
LANDER	0	0	**0**
CONSERT	90	91	**181**
RSI	0	0	**0**
Total instrument allocation	68830	63138	**131968**

Table 4. Instrument Science data allocation. *Typical science data volume allocation per instrument, for the April 2015 period (MTP015). The MTP values are sub-divided into two segments of 2 weeks each, to allow more flexibility in the allocation of the resources. The total allocated (131.9 Gb) is slightly higher than the available (131.5Gbit).*

C. Data volume during the instruments command iterations cycle

The Operations and Planning group generates, from the harmonized observation plan, a initial set of monthly operation timelines and spacecraft pointings that the teams will use as templates to create their command requests. The initial allocated data volume is reproduced in these draft command files. The iterations are held weekly, lasting ~6 weeks.

The data volume is defined and computed at the Science Ground Segment in two ways:

- Instrument modeling. The data volume is computed by the Experiment Planning Sytem (EPS) using a modeling defined in the Experiment Description Files (EDF). The data generation is triggered by the requested instrument commands and parameters in the operational payload operation requests. Examples of this method are the navigation images, Rosina and RPC.
    - The NAVCAM data volume is triggered by the image acquisition sequence, and the sequence parameters models the size of the image (which columns and rows are requested from the CCD readout).
    - The RPC instrument estimates the data volume generation as a constant profile that depends on the operation mode. The operational sequences makes the logic transitions between the different instrument modes and module states, that has defiend different bit rates. The next example extracted from the real RPC EDF illustrates the module states (OFF, Quiet, Min, Normal and Burst) of the subunit Langmuir Probe Instrument (RPC-LAP).

```
Experiment: RPC
  Module: LAP
      Module_state: OFF
          MS_power: 0 [Watts]
          MS_data_rate: 0 [bits/sec]
      Module_state: QUIET
          MS_power: 2.07 [Watts]
```

```
    MS_data_rate: 0 [bits/sec]
Module_state: MIN
    MS_power: 2.07 [Watts]
    MS_data_rate: 1.6 [bits/sec]
Module_state: NORMAL
    MS_power: 2.07 [Watts]
    MS_data_rate: 62.5 [bits/sec]
Module_state: BURST
    MS_power: 2.07 [Watts]
    MS_data_rate: 2253 [bits/sec]
```

- Z-record. The data volume is extracted by the Experiment Planning Sytem from an explicit data statement (Z-record) in the instrument sequences. Each z-record specifies a data value that is constant and valid from a specified time (given as a delta time) until the next specified value. Each sequence can contain a number of z-records, allowing for a customized resource profile. An example based on a real Virtis sequence is provided next. The Virtis-M (mapper) sequence AVRF010A is declared with parameters (VVRG0051 ...) and a z-record DATA_RATE_PROFILE with 3 entries.

```
13-Jan-2016_09:47:00 AVRF010A ( \ # VIRTIS M Sci SSMM v10
             VVRG0051 = 1              \       # M_ERT
             VVRG0053 = 5              \       # M_ACQ
             VVRG0035 = 160            \       # M_CCD_EXPO
             VVRG0037 = 25247          \       # M_ALPHA_FIRST
             VVRG0038 = 62847          \       #M_ALPHA_LAST
             VVRG0041 = 40             \       # M_DBCK_RATE
             VVRG0048 = 255            \       # M_IR_DET_OFF
             DATA_RATE_PROFILE =  \
             00:00:00   34124 [bits/sec]\    # [bits/sec]
             00:25:00     1334 [bits/sec]\    # [bits/sec]
             00:30:00          0  [bits/sec])    # [bits/sec]
```

The data volume is communicated to the Flight Control Team for final verification before upload only with Z-records (not modeling). One of the key aspects of the validation of the commanding files is that the teams should stay within the allocated volumes, and that peaks in production should not result in memory overflows. In case of conflicts, the RSGS check software (Mission Analysis and Payload Planning System, MAPPS[4]) will generate warnings and the operations engineer will instruct the affected teams for modifications of their operation requests. Figure 4 shows a typical SSMM utilization graph, used during the iterations cycle to check visually the status of the data stores.

D. Data downlink capacity changes during the nominal planning cycle

The predicted data downlink usually changes during the detailed nominal planning iterations. First, the ground station allocations can be modified due to request or agreements with other missions. Second, the available downlink capacity is affected by the pointing. Rosetta performs its dumps via the the High Gain Antenna (HGA), a steerable antenna that can point to Earth independently of the payload field of views direction. But sometimes the pointing direction or slew characteristics exceed the HGA tracking capability, and the HGA cannot keep continuous Earth tracking: a communications outage happens. The most usual occurrence is the so called HGA flip, where the antenna turns 180 degrees to be able to keep a tracking pattern. On fewer occasions the tracking direction enters in a small but forbidden area where the antenna mechanism cannot point to, due to spacecraft body and appendices blockage. These outage times can be computed and introduced only at a late stage, when the pointing is known. In this case, the total downlink capacity needs to be re-computed and in some critical cases the attitude has to be modified to avoid significant data losses.

The total downlink capacity is re-analysed weekly for each planning cycle step, to identify significant modifications from the initial estimation. If there is enough time to re-iterate, a data volume reallocation might be done, in coordination with the Liasion Group and the experiment team.

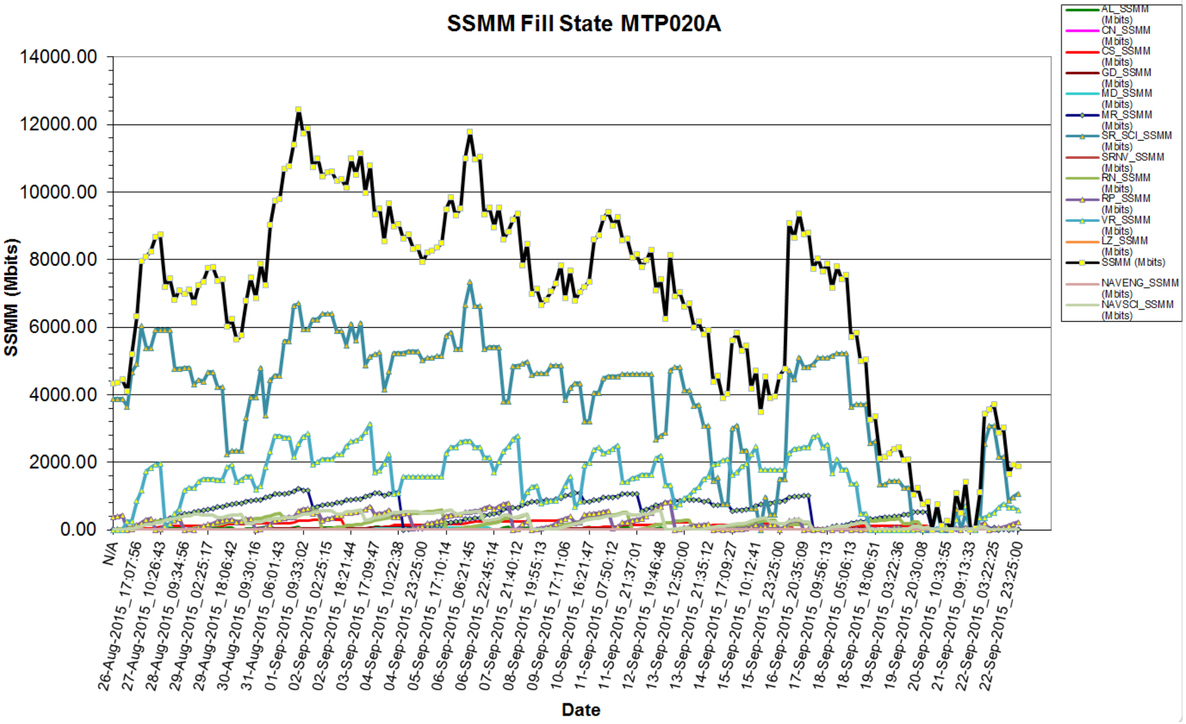

Figure 4. Rosetta SSMM fill state graph. *The graph shows the different instrument SSMM fill status during MTP020 (Aug-Sep 2015). There is an initial 4.4Gbit carry over from the previous period. The first 3 weeks are periods of high production rate, and the SSMM never gets fully empty until around 19th Sep.*

E. Downlink optimization using priority dump

Rosetta also has the capability to select priorities for each instrument during each ground station pass. RSGS is responsible for the data management and uses an optimization tool called DALLOC [5,6]. The DALLOC is a software tool that optimizes the data downlink during a given period, usually an MTP, balancing the individual instrument SSMM filling loads. It avoids overflows on any instrument due to the particular instantaneous production rate and tries to reduce the data latency for the requested experiments. It was developed by the Artificial Intelligence Group at NASA/JPL with RSGS requirements. It is worth noting that the Rosetta downlink problem is particularly difficult as not only are there a large number of buffers and events, additionally downlinks and data production events (observations) have significant temporal extent. As indicated above downlinks over 16 hours in duration and observations of days or weeks in duration are frequent occurences. This necessitated more general data management models than previously used that presumed that data generation events and downlinks do not overlap [7,8].

Once the instrument commands are verified by the Operations and Planning group, and the average data volume productions are according to the allocated values, the DALLOC optimization is performed. Two steps are needed.

First, once the planning period configuration and instrument payload requests are in place, MAPPS will be run over the whole MTP, with no downlink optimization yet used. Therefore instantaneous overflows might happen. At this stage, both the allocated data generation and the data downlink policy for the MTP are known and fixed. DALLOC will analyse this information and come up with optimized downlinks, which are fed back into MAPPS. The end result should be a conflict free, balanced downlink.

A second step is needed to account for the modifications of the data volume profile. These changes come every week with the reception of new or updated weekly-span requests from the instruments. It might also be needed when there are changes in the ground station passes. In some cases a data volume reallocation exercise is needed, such as sudden reductions of ground station availability, or for general data volume under/overestimation for the MTP. A MAPPS re-run over the whole MTP is used to provide the new conditions for DALLOC optimization, and the resulting timeline file overrides the existing one.

DALLOC generates a timeline file containing a set of payload downlink priorities for each available dump pass. These are sent to RMOC with the rest of instrument commands, for upload to the spacecraft.

F. Data latency management

Data latency is the time a packet takes to reach ground after it was generated on-board. Several Rosetta teams need the data on ground within a few hours/days from acquisition, so that they can feed the results back into the upcoming observations. This is the case, for example, of the COSIMA built in microscope camera (images of cometary dust grains) and MIDAS atomic force microscope (context images are quickly dumped to identify candidate particles for re-scans at high resolution). Other requests arrive during infrequent situations where the instruments need software patches, or perform interactive checks of new functionalities.

The RSGS software MAPPS [4] computes the data latency. Latency is a secondary constraint to the downlink optimizer tools, so if a team request is not met, often a manual intervention into the priority dump file is followed. If the request is more generic (like Cosima, Midas, and Giada) the DALLOC tool can be set to maximize the data dump priority for these teams. This is only recommended for such teams with low data volume generation.

G. Instrument data reallocations

After the teams have delivered their operations request during the planning cycle iterations, it might happen that some of the data volume is finally not fully utilized by a team. In these cases it is possible that a reallocation of data volume is performed. If there is enough time to re-iterate, a data volume reallocation exercise is done, in coordination with the Liasion Group and the experiment team. In case that for a significant period (days) the ground stations are heavily underused, and the instrument teams cannot increase anymore their data generation, the RSGS might release some of the booked stations to benefit other missions.

H. Modifications during ground pass execution

There might be live technical problems that prevent the actual dumps in the ground stations. Some examples are false signal locks, station unavailability, and halts due to extreme weather conditions (snow, wind). In some occacions there are full pass cancellations due to emergency actuations to other spacecraft. These usually rare problems produce deviations in the actual SSMM onboard states from the predictions, invalidating the current priority dumps and raising the risk of overflows and data loss.

I. Special payload operations

Occasionally the instrument teams request an extraordinary payload operation, commanded by means of a Payload Direct Operations Request (PDOR). These type of operations are handled directly between the science team and the Mission Operations Center, and it will most likely cause an increment of data volume which is not modeled in the nominal commands handled by the RSGS. This is the case for example of last minute instrument health actuations, corrections of wrong instrument settings, and last minute target of opportunities known just before execution. The Operations and Planing group, in this case, procedurally creates some pure modeling sequences to account for the extra data volume requested, but ensuring that no extra command will be generated for upload. Usually the effect of this anaylis is a reconsideration of the data volume allocation to the teams, and or a change to the priorities schema to re-balance correctly the load on the different instrument data stores.

J. Post-execution analysis and feedback into the planning cycles

The RSGS performs a post-execution analysis of the actual generated and dumped data[9]. It consists of two distinct checks.

The first one is the instrument data generation trend. RSGS has deleloped a tool that monitors weekly the average data dump for each instrument, and compares it to the predictions. The goal is to identify discrepancies in data volume generation, and eventually fine tune the production models of the instruments. In case of periodic discrepancies on several consecutive planning periods, the affected team is contacted to either explain it (for temporary features like over/under compression when the imaged area is too dark/ bright) or to change the modeling of the data volume generation. This has helped on several occasions to refine the experiment models.

The second check is the actual SSMM fill status at any given time. Unfortunately there is no mechanism nor procedure to directly know the actual Rosetta SSMM filling. The RSGS has therefore developed an indirect onboard memory estimator, following this rationale. The Science Ground Segment knows the last instrument data dumped to ground, following the instrument identificator and time stamp in the telemetry packet header. The onboard SSMM works in a FIFO scheme (First In First Out) and therefore the stored data is sequential in time. By compensating in the predictions all the data that was dumped at the time of arrival of the last packet, a much better prediction using the future instrument generation can be achieved. This tool has been used on several occasions, especially in critical periods with high data volume fillings. In this case the RSGS will reset the predictions in the modeling tools, to be able to accommodate an increased data generation for the following days. Please note that usually the SSMMs are slightly less occupied than predicted, because of the built-in margins in the ground station dump times and the use of worst case values for housekeeping.

IV. Future improvements in dump strategy

For the last months of the Rosetta mission, the data downlink capacity will be severily limited, as the distance to Earth will increase to more than 4 AU. To mitigate the low bit rates, new ways to dump data are being implemented between the RMOC and the RSGS, based on the existing established working procedures. In particular the so called 'blind dumping' will be used. In this new Rosetta downlink schema, some instruments will be entitled to generate an extra leap of data and dump it to Earth before the ground station has confirmed the nominal start of locked communications. The dump duration will increase by around two times the OWLT. The risk associated with losing data due to initial false locks or ground station unavailability will not be zero, but it will be surpassed by the overall data boost.

V. Conclusion

The Rosetta data volume statistics are quite impressive. From early 2014 until the time of writing (March 2016), Rosetta has downlinked to Earth ~2.8 Tbit of data, using around 10.000 hours of ESA ESTRACK ground stations, 2500 hours of NASA DSN 35m stations, and 2400 hours of NASA 70m antennas. Despite the enormous amount of data received, and the complexity to harmonize the data generation of the capable payload suite (11 instruments on the Orbiter and 10 on the Lander), the RSGS has achieved a well balanced SSMM data stores utilization such that no data has been lost due to RSGS planning or scheduling problems. This is a significant achievement for such a complex mission.

In addition, the robust planning process has allowed to perform last minute modifications to the experiment observation plans, to cope with the highly dynamic cometary environment and the complex comet navigation and pointing. This required a large multi-team effort to update on time the commands, account for the extra data volume, and to rebalance correctly the SSMM stores filling and assigned downlink priorities. But it proved successful.

Appendix

ROSETTA ORBITER SCIENCE INSTRUMENTS		PHILAE LANDER SCIENCE INSTRUMENTS	
Alice	Ultraviolet Imaging Spectrometer	APXS	Alpha Proton X-ray Spectrometer
CONSERT	Comet Nucleus Sounding Experiment by Radio wave Transmission	ÇIVA and ROLIS	Panoramic and microscopic imaging system
COSIMA	Cometary Secondary Ion Mass Analyser	CONSERT	Comet Nucleus Sounding Experiment by Radio wave Transmission
GIADA	Grain Impact Analyser and Dust Accumulator	COSAC	Cometary Sampling and Composition experiment
MIDAS	Micro-Imaging Dust Analysis System	PTOLEMY	Evolved Gas Analyser
MIRO	Microwave Instrument for the Rosetta Orbiter	MUPUS	Multi-Purpose Sensor for Surface and Subsurface Science
OSIRIS	Optical, Spectroscopic, and Infrared Remote Imaging System	ROMAP	Rosetta lander Magnetometer and Plasma Monitor
ROSINA	Rosetta Orbiter Spectrometer for Ion and Neutral Analysis	SD2	Sample and Distribution Device
RPC	Rosetta Plasma Consortium	SESAME	Surface Electric Sounding and Acoustic Monitoring Experiment
RSI	Radio Science Investigation	APXS	Alpha Proton X-ray Spectrometer
VIRTIS	Visible and Infrared Thermal Imaging Spectrometer		

Table 5. List of Rosetta Orbiter and Lander Instrument packages.

Acknowledgments

The effort required to plan the Rosetta science observations, to implement the detailed pointing and instrument timelines, and to execute the onboard and ground station commands has been considerable during several years, especially after Rosetta started the close operations at the comet in spring 2014. The data volume management, as described in this paper, was a significant part of this process. Therefore the authors want to thank the main teams involved in the process, RMOC, RSGS and the PI teams, for their dedication, proffesionality and in the end successful results in this unique mission.

References

[1] Pérez-Ayúcar, M., Ashman, M., Almeida, M., Costa, M., García Beteta, J.J., Hoofs, R., Marín Jaseli, J., Merritt, D., Nespoli, F., and Sánchez, E., "The Rosetta Science Operations and Planning Implementation", *Acta Astronautica* (submitted for publication).

[2] Vallat, C., et al., "The science planning process on the Rosetta mission", *Acta Astronautica* (submitted for publication).

[3] Hubault, A., and Urbanek, J., "Rosetta - Operations Rules and Constraints Document", *ESA Technical Note* RO-ESC-RS-5301

[4] Plas, P. van der, García-Gutiérrez, B., Nespoli, F., Pérez-Ayúcar, M., "MAPPS: A science planning tool supporting the ESA Solar System missions" *Spaceops 2016 Conference*, AIAA, Daejeon, South Korea (submitted for publication)

[5] Rabideau, G., Nespoli, F., and Chien, S., "Heuristic Scheduling of Space Mission Downlinks: A Case Study from he Rosetta Mission", *International Workshop on Planning and Scheduling for Space*, Buenos Aires, Argentina, July 2015.

[6] Chien, S., Rabideau, G., Tran, D., Troesch, M., Doubleday, J., Nespoli, F., Perez-Ayucar, M., Costa-Sitje, M., Vallat, C., Geiger, B., Altobelli, N., Fernandez, M., Vallejo, F., Andres, R., Kueppers, M., Activity-based Schedulilng of Science Campaigns for the Rosetta Orbiter, Invited Talk, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2015)*. Buenos Aires, Argentina, July 2015.

[7] A. Cesta, G. Cortellessa, S. Fratini, A. Oddi, and N. Policella., "An Innovative Product for Space Mission Planning: An A Posteriori Evaluation" *Proceedings International Conference on Automated Planning and Scheduling*, pp. 57-64, 2007, AAAI Press.

[8] Righini G, Tresoldi E., "A mathematical programming solution to the Mars Express memory dumping problem", *Systems, Man, and Cybernetics, Part C: Applications and Reviews*, IEEE Transactions on. 2010 May;40(3):268-77.

[9] O'Rourke, L., Cuevas, M.A, Aberasturi, M., Villacorta, A., Pesciullesi, P., Diaz, J., Kueppers, M., Taylor, M., Martin, P., Jansen, F., "Rosetta Science Ground Segment Observation Quick Look System", *Acta Astronautica* (submitted for publication).