Using an Artificial Intelligence Tool to Perform Science Data Downlink Planning as Part of the Mission Planning Activities of Mars Express

Erhard Rabenau
NOVA Space Associates Ltd
11 Kingsmead Square, Bath, BA1 2AB, UK
Erhard.Rabenau@esa.int

Amedeo Cesta, Angelo Oddi
ISTC-CNR, Institute for Cognitive Science and Technology
Via S. Martino della Battaglia, I-00185 Rome, Italy
Amedeo.Cesta@istc.cnr.it, A.Oddi@istc.cnr.it

Alessandro Donati, Thomas Léauté
European Space Operations Centre
Robert-Bosch-Str. 5, 64293 Darmstadt, Germany
Alessandro.Donati@esa.int, Thomas.Leaute@esa.int

Abstract

Solving the science data downlink problem for Mars Express (MEX) has been a challenge. An artificial intelligence (AI) based tool, MEXAR2, has been integrated in the mission planning process that allows the user to quickly and effortlessly generate and analyse a dump plan for any planning period. MEXAR2 is the first AI-based tool in an operational environment on-ground at ESOC and it is the successful implementation of a prototyping study performed before the launch of MEX. The paper explores the challenging requirements levied on the MEX science data downlink, the different solutions for generating a dump plan before implementing MEXAR2. It further describes in detail the interfaces and functions of the software and provides an evaluation of the operational benefits from using the application compared to previous solutions.

Nomenclature

CCSDS	=	Consultative Committee for Space Data Systems
DSN	=	Deep Space Network
ESA	=	European Space Agency
Gbit	=	1000 million bits
kbps	=	kilobits per second
Mbit	=	1 million bits
MDB	=	Max Dump Problem
MEX	=	Mars Express
MEXAR	=	Mars Express scheduling Architecture
MMI	=	Man Machine Interface
MPS	=	Mission Planning System
SSMM	=	Solid State Mass Memory

Introduction

Mars Express (MEX) is ESA's first mission to another planet. Even though classed as a 'small' mission, the mission objectives are manifold and include the exploration of Mars with high resolution stereo imaging and mineralogical mapping of the surface of the planet, subsurface radar sounding, atmospheric sounding as well as environmental and radio studies. This complex mission operations concept has provided operational challenges in many fields, one of them being the downlink of the collected science and housekeeping data. After arriving at Mars in late 2003, MEX has been generating more science data than originally planned for. This was mainly due to the introduction of additional ground stations thus increasing the available downlink time and resulting volume. MEX is currently in the first mission extension which will last till late 2007. The baseline mission design committed to downloading 500 Mbits of science data per day. The actually achieved downlink volume is on the order of several Gbits per day depending on the downlink bit rate and time available. In the MEX mission planning concept, downlink time and therefore volume of data generated is a controllable resource, just as power. The paper presents the challenges involved with generating a downlink plan and the resulting solution, which is another first for an ESA mission: the implementation of an artificial intelligence software application which has been tightly integrated in the overall mission planning process.

The MEX Downlink Problem

The following section describes the MEX downlink problem in detail.

MEX operations are performed in a store-and-forward manner, i.e. all telecommands, science and housekeeping data are stored in a Solid-State Mass Memory (SSMM) unit till execution or downloading to earth (also referred to as data dumping). The SSMM comprises dedicated packet stores (which can be perceived as data directories or memory areas) for each instrument and housekeeping data.

Science data is written to so-called cyclic packet stores which can be perceived as circular buffers. When more data is written to the store than its defined capacity, the oldest data is overwritten. Data from cyclic packet stores can be dumped by opening the file for reading and closing it after a time interval, corresponding to the dump duration. The management of the read and write pointers is automatically performed by the SSMM operating system. The size of the science packet stores has been designed to accommodate about 2 days of typical science data generation. They range from 300 Mbits to 4 Gbits in volume. Each instrument has been allocated one science packet store, with the exception of the camera, which has been allocated two.

Housekeeping data is stored in so-called non-cyclic packet stores. They are tailored for managing large numbers of small packets, e.g. events, telecommand acknowledgements, etc. Non-cyclic packet stores are dumped in their entirety, i.e. as defined by the packet store size. They can be likened to a bucket. Once the container is full, any additional data will be lost, i.e. the newest data will be lost. Non-cyclic packet stores range from 5 Mbit to 200 Mbits in volume. Housekeeping packet stores comprise an A and B store which are swapped every day at midnight, after which the data from the previous day is dumped to the ground (this is done to facilitate data management on board and on the ground). Packet stores, either cyclic or non-cyclic, can only be dumped in sequence, one at a time. Parallel dump operations are prohibited for SSMM performance reasons. Housekeeping dumps have a higher priority than science and need to be downloaded as soon as possible after a packet store swap.

Science operations are planned at medium-term planning level and have to take into account the available downlink windows, the downlink bit rate, the instrument data generation rate and the available packet store volume, the main driver being that all generated data should be dumped before the packet store contents gets overwritten (this constraint evolved from a more stringent one that all generated data should be dumped within a day – this was mostly due to the processing restrictions for generating a dump plan). As instrument operations and downlink are a controllable resource, ground station passes and instrument operations will be traded off for maximum science return. Therefore, ground station passes will not only be interrupted by fixed spacecraft maintenance slots, occultations and eclipses (during which the transmitter will be switched off), but also by science operations. As a result, a ground station pass may be split into several downlink opportunities or windows ranging from less than 30 minutes to several hours (any downlink opportunity above a certain threshold will be taken in order to maximise the downlink volume). The resulting instrument operation timeline typically contains 1 or 2 instrument operations per orbit, or up to 8 instrument operations per day, generating between 80 Mbit and several Gbits. While the observing instruments are switched off after each observation, the ASPERA instrument is constantly producing data over a whole day. The accurate prediction of the amount of data generated by an instrument is of utmost importance for generating a dump plan. A scheme was therefore introduced to specify the data rate as a profile over time in the standardised payload operations requests that are generated by science planning and the instrument science teams. For some instruments the actual amount of data generated may vary from the predicted values by 10% or more depending on the applied compression algorithm of the processing software. Depending on the distance of Mars from Earth, the downlink bit rate varies from 28 kbps to 182 kbps. The downlink bit rate may vary between downlink windows due to the allocated antenna size. During the course of the nominal mission, science planning very often fully utilised the available downlink channel leaving no or very little margin for retrieving data in case of data loss during downlink.

In summary, the following constraints have been applied in order to achieve a feasible dump plan:

- The available downlink windows are grouped in a so-called dump day which contains all downlink windows of a given day starting with the first downlink window after midnight;
- Each downlink window is allocated a dedicated downlink bit rate;
- Housekeeping telemetry is to be dumped as early as possible in a dump day;
- The generation of a second packet store for the camera is not modelled in the data profile of the operations request, as the standardised interface does not support multiple packet stores per instrument.

Solutions

MEXAR Study

The downlink of MEX science data has always been perceived as a challenge. Before launch, a study was performed that addressed the downlink problem (Cesta, Oddi, Cortellessa and Policella 2002). It addressed a solution by application of artificial intelligence technology, referred to as the MEX memory dump problem (MDP). But because of incomplete operational requirements and missing interface descriptions, the study, even though it provided interesting results, did not progress beyond the final presentation.

Manual Operations

During cruise and Mars in-orbit commissioning, all dump operations were scheduled manually. This was, however, quickly abandoned because of the workload involved and the unlimited potential for error.

Dump Generation Tool

Next, a tool was developed in Visual Basic that semi-automated the generation of a dump plan. The tool used as inputs:

- A dump windows file generated by mission planning at ESOC as part of the medium-term planning process taking the latest ground station allocations and science pointing timeline into account;
- Instrument data share allocations, i.e. data volumes to be dumped in a dump day on an instrument-by-instrument basis.

The tool allowed a quick generation of a dump plan but was limited: data share allocations were not linked with the data generation time and the distribution of dumps over a day or several days proved to be difficult. As the science data generation increased, it became more and more important to dump science data as early as possible after generation in order to avoid overwrites of the packet store. Several iterations of manually changing the data share allocations were usually necessary to arrive to a feasible dump plan.

The output of the tool were dump operations requests as inputs to the MEX Mission Planning System (MPS), a dump plan and a dump list with which the packet store contents could be monitored in an MSEXCEL spreadsheet after merging the data volumes from the operations requests. As the application was PC-based, input and output files had to be transferred to the MPS across a firewall.

MEXAR2

MEXAR2 was born during an IWPSS at ESOC when the MEXAR study team proposed to implement an operational prototype using the software application developed during the study. Based on the experience with the dump generation tool, a set of requirements and operational data could be generated within a very short time. The assumption was made that the interfaces as defined for the dump generation tool should be maintained in order to allow a seamless transition between the two systems. The most import change was that MEXAR2 should be able to read the data profiles in the payload operations requests thus establishing the generation time of the science data.

A first prototype was quickly available and, after some tests and bug fixes, could be used in the operational mission planning environment, already providing an improved mechanism compared to the previous approach. Several prototype deliveries with added functionality followed. Each delivery was tested by the mission planning team and anomaly reports were generated. The MEXAR2 development team quickly reacted with bug fixes, such that the improved application could quickly be used operationally. The final version of the tool runs within a user-friendly man-machine-interface (MMI). It has been developed in JAVA and therefore runs on PC and UNIX platforms (see Figure 2).

The MEXAR2 dump planning tool has been tightly integrated with the mission planning tools at ESOC (see Figure 1).

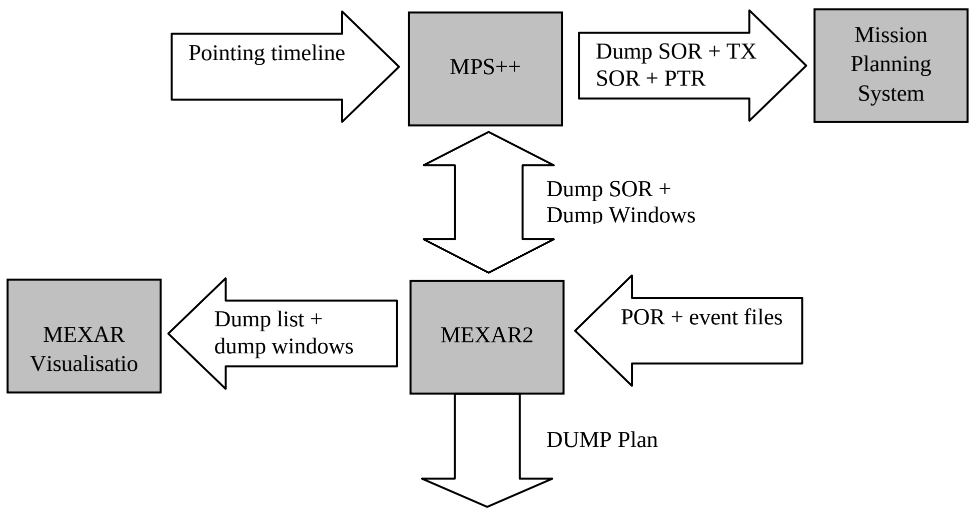

Figure 1: Integration of MEXAR2 with the mission planning tools at

It requires the following inputs:

Configuration File Entries. The configuration file can be created from the MMI. Different configuration files can be created and can be selected as required for a planning session. A configuration contains:

- A dump windows file as automatically generated during the medium-term data planning process, organised in dump days with the downlink bit rate defined for each window (a dump windows file is generated for the period of a medium-term plan – usually 28 days – with an average of 4 ground station passes per day);
- Payload operations requests containing the data rate profile over time (a medium term may contain over 400 payload operations requests with around 10 profile entries for some instruments);
- An orbit event file to determine the time of the data rate profiles (as times in the payload operations requests are

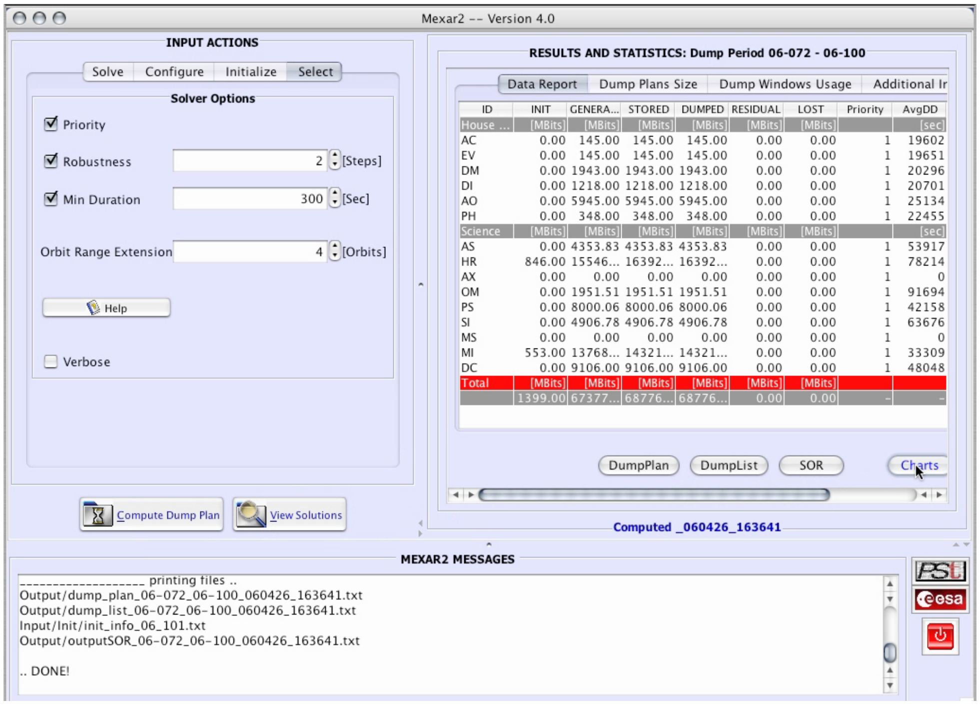

Figure 2: MEXAR2 User Interface

specified relative to an orbit event – usually the pericentre crossing time);

- A domain specification file (names and volumes of the packet stores defined for MEX, priority definitions, robustness settings, packet swap command sequences, delay times);
- MEXAR2 event files:
  - o Absolute events: allow to specify packet store dumps for a specific dump day;
  - o Periodic events: allow to specify packet store dumps repeatedly on each dump day;
  - o Synchro events: allow to specify packet store dumps that are synchronised with the generation of data in another packet store, e.g. each time the camera performs an observation and generates data for the science packet store, an entry is generated in the auxiliary packet store.

User Inputs from the MMI. The user has to specify several parameters before executing the software. These are:

- Dump day range;
- Selected configuration;
- Initial values (residual data at the start of the dump plan) in the packet stores and the packet store swap identification (if MEXAR2 is used in a continuous dump plan generation mode, i.e. with contiguous dump days, then the initial values will be taken from the previous session);
- Whether priority handling is to be executed (if priorities are defined in the domain specification file);
- Level of robustness: robustness is a feature that tries to distribute the packet store dumps in order to avoid packet store overwrites;
- Minimum duration of a data dump;
- Payload operations request overlap: this feature takes the specified number of orbits before the start of the plan into account in order to avoid loss of data, e.g. for instruments that are already producing data at the start of the plan.

The outputs of the MEXAR2 application are:

Dump operations request template. This template defines the command sequences and absolute execution times. It is compatible with the input required to the operations request generation application, in the mission planning software that generates a spacecraft operations request from the template, taking the parameter settings of the operational database into account.

Dump plan. This is a list of all the dump activities in a time-ordered fashion per dump day, downlink bit rate, volume to be dumped, volume dumped, remaining volume and time taken to dump.

Dump list. This is a list of the data volume profiles and dumps generated in a time-ordered fashion as input to other applications, e.g. spreadsheets, graphical visualisation tools.

Statistics file. This is an overview of the overall data downlink capacity, number of dump plan entries per dump day, overall data volume generated, dumped, residual and lost data due to overwrites.

Unused dump windows file. This provides information about the usage of the dump windows and lists the unused portions of a dump window that can be used by the operations team to recover data lost during a downlink.

MEXAR2 allows the user to quickly perform what-if analyses for dumps in the same planning range by generating and comparing dump plans with different configuration settings, e.g. to determine the impact of a robustness setting.

Additional tools have been developed by a third party to visualise the dump plan. This JAVA-based tool has been integrated in MEXAR2.

Conclusion

Space Missions are challenging domains for applying Artificial Intelligence techniques for planning and scheduling. MEXAR2 is the operational implementation of such an on-ground system featuring an automated problem solving technique. The system performs the operational tasks reliably and in a time-efficient manner, providing an easy-to-use and intuitive user interface. MEXAR2 provides science data on ground as early and complete as possible after generation (this is not a mission requirement but an intrinsic feature of the used max dump algorithm) and allows for the minimisation of data losses due to overwrites. The use of MEXAR2 by the MEX mission planning team offers an overall time saving of about 50% in the dump planning cycle compared to the previous implementation due to a comprehensible end-to-end view of the dump activities including the add-on of a third party visualisation software.

As it uses standardised interfaces, it could be easily adapted for other missions in the end-to-end dump planning activities of the mission planning team at ESOC.

As a direct result of the success of MEXAR2, a similar application for solving the MEX uplink problem is being developed.

References

Cesta, A., Oddi, A., Cortellessa, G., Policella, N. 2002. Automating the Generation of Downlink Spacecraft Operation in Mars Express: Analysis, Algorithms and an Interactive Solution Aid, Technical Report, MEXAR-TR-02-10, ISTC-CNR [PST], Italian National Research Council.

Oddi, A.,and Policella, N. 2004. A Max-Flow Approach for Improving Robustness in a Spacecraft Downlink Schedule. *Fourth International Workshop on Planning and Scheduling for Space - IWPSS '04* pp. 151-158.