

Please select category below:

Normal Paper

Student Paper

Young Engineer Paper

Navigation Challenges during ExoMars Trace Gas Orbiter Aerobraking Campaign

Gabriele Bellei ¹, Francesco Castellini ², Frank Budnik ³ and Robert Guilanyà Jané ⁴

¹ DEIMOS Space located at ESA/ESOC, Robert-Bosch-Str. 5, Darmstadt, 64293, Germany

² Telespazio VEGA located at ESA/ESOC, Robert-Bosch-Str. 5, Darmstadt, 64293, Germany

³ ESA/ESOC, Robert-Bosch-Str. 5, Darmstadt, 64293, Germany

⁴ GMV INSYEN located at ESA/ESOC, Robert-Bosch-Str. 5, Darmstadt, 64293, Germany

Abstract

The ExoMars Trace Gas Orbiter satellite spent one year in aerobraking operations at Mars, lowering its orbit period from one sol to about two hours. This delicate phase challenged the operations team and in particular the navigation system due to the highly unpredictable Mars atmosphere, which imposed almost continuous monitoring, navigation and re-planning activities. An aerobraking navigation concept was, for the first time at ESA, designed, implemented and validated on-ground and in-flight, based on radiometric tracking data and complemented by information extracted from spacecraft telemetry. The aerobraking operations were successfully completed, on time and without major difficulties, thanks to the simplicity and robustness of the selected approach. This paper describes the navigation concept, presents a recollection of the main in-flight results and gives a retrospective of the main lessons learnt during this activity.

Keywords: ExoMars, Trace Gas Orbiter, aerobraking, navigation, orbit determination, Mars atmosphere, accelerometer

Introduction

The ExoMars program is a cooperation between the European Space Agency (ESA) and Roscosmos for the robotic exploration of the red planet. The first ExoMars mission comprised an orbiter, called Trace Gas Orbiter (TGO), and an Entry, descent and landing Demonstrator Module (EDM) [1, 2]. It was launched by a Proton-M rocket from Baikonur into a direct transfer to Mars on 14th March 2016. After a successful Mars Orbit Insertion manoeuvre on 19th October the TGO spacecraft entered a 4-sol highly elliptical orbit which was changed, through a series of chemical manoeuvres, into a 1-sol, 74 degrees inclination orbit [3]. From here, the spacecraft started aerobraking operations on 15th March 2017, with the purpose of reducing the orbital period down to about 2 hours with limited use of propellant. The aerobraking campaign was successfully completed on 20th February 2018, after 952 transits through the Martian atmosphere for a total ΔV of about 1.02 km/s, and started science operations in mid-April 2018, after acquiring the operational orbit via a series of chemical correction manoeuvres.

The TGO aerobraking represented a major challenge for all the parties involved, in particular the project team, the industrial consortium and the mission control team, due to its inherent complexity, extended critical operations and associated risks. A broad description of TGO

aerobraking design and operations was provided in Ref. 4, while a more detailed discussion on the Flight Dynamics aspects can be found in Ref. 5.

The purpose of this paper is to provide insight about some specific topics of TGO aerobraking, in particular those concerning the navigation operations carried out by the Flight Dynamics team at the ESA's European Space Operations Centre (ESOC). In the following sections, a brief summary on the aerobraking technique and the requirements it imposes on the navigation system will be provided, followed by a description of the navigation concept and a report on the main results obtained during in-flight operations. Finally, a recollection of the main lessons learnt during this phase will be given.

Aerobraking at Mars

The aerobraking technique consists in the use of the atmospheric drag to reduce the energy of a spacecraft and lower its orbit period in a controlled manner. It is typically used to reach quasi-circular orbits starting from highly elliptical orbits, following a planetary insertion manoeuvre, and permits propellant savings at the cost of an additional complexity in satellite design and operations. For a typical low Mars orbit satellite, the ΔV savings are in the order of 1 km/s. The technique was pioneered by the NASA mission Magellan around Venus in 1993 [6] and since then successfully used by the Jet Propulsion Laboratory (JPL) for Mars Global Surveyor [7, 8], Mars Odyssey [9, 10] and Mars Reconnaissance Orbiter [11, 12]. Although a very basic aerobraking experiment was carried out at the end of the ESA mission Venus Express [13], TGO represented the first ESA satellite including this technique in the baseline mission definition, requiring the complete design of all the necessary space and ground segment systems.

The Martian atmosphere at the typical braking altitude behaves as a flow at a free molecular or transition regime [14]. A satellite passing through its upper layers experiences a free-flow drag force proportional to the square of the incoming flow relative velocity. The energy dissipation will produce a heat flux which is proportional to the cube of the incoming flow velocity. The effect of a full atmospheric pass is a reduction in the spacecraft velocity (ΔV) and kinetic energy and an increase in the temperature of the surfaces exposed to the flow, due to the integrated effect of the heat flux (heat load).

Given the atmospheric density ρ at a given position, and the magnitude of the relative incoming velocity of the atmospheric particles V_{in} , the mathematic expressions are provided for the free flow *dynamic pressure* (Eqn. 1), *heat flux* (Eqn. 2) and *heat load* (Eqn. 3, integrated over a single atmospheric pass). All the quantities are intended per unit surface.

$$p_{dyn} = \frac{1}{2} \rho V_{in}^2 \quad (1)$$

$$\dot{q} = \frac{1}{2} \rho V_{in}^3 \quad (2)$$

$$Q = \int_{AB} \dot{q} dt \quad (3)$$

The aerobraking design should try to maximise the effective ΔV resulting from each pass, thus reducing the time needed to reach the target orbit period, while preserving the safety and integrity of the satellite and its instruments. An *aerobraking corridor* is defined, having as the upper limit any of the structural and thermal limits of the satellite and as the lower limit the minimum ΔV that will guarantee the completion of the aerobraking within the time allocated. The pericentre height must be frequently adjusted by means of propelled manoeuvres, in order to maintain the aerobraking within the corridor, which is a challenging task due to fast spatial

and temporal variations of the atmosphere density that are difficult to predict. These manoeuvres, executed at apocentre and controlling the pericentre height, will be called in the following Pericentre Raising Manoeuvres (PRM) and Pericentre Lowering Manoeuvres (PLM). Aerobraking can be conventionally divided into three main phases: a *walk-in* phase, during which the pericentre altitude is lowered gradually with progressive manoeuvres to characterise the current status of the Martian atmosphere; a *main* phase in which the height is continuously adjusted to keep the satellite within the predefined corridor and most of the altitude and period reduction is achieved; and finally a so called *end-game* phase, during which the orbit period is low, the satellite spends a considerable fraction of its orbit within the atmosphere and the efficiency of a single braking manoeuvre is much reduced in terms of orbit period reduction. The aerobraking is terminated by a *walk-out* manoeuvre raising the pericentre above the detectable atmosphere.

The Aerobraking Navigation Problem

TGO Spacecraft Design and Navigation Requirements

The TGO satellite was designed already with aerobraking operations in mind. Its Guidance, Navigation and Control (GNC) architecture includes a dedicated aerobraking mode [15, 16], encompassing the necessary on-board autonomy to deal with the uncertainty of the dynamic environment, in particular the variability of the Martian atmospheric density that can affect both the aerodynamic loads and the orbit prediction accuracy. However some of TGO features, summarised in Table 1, were driven by other aspects, including its role as a carrier for the EDM and the partial reuse of an existing platform, and they are therefore not ideal for aerobraking. Even after the EDM was released and most of the fuel was consumed during Mars orbit insertion and orbit change manoeuvres, the satellite is heavy, with a very high ballistic coefficient. Its agility is limited by the relatively small reaction wheels and a fairly inefficient force-free thruster configuration, which increases the 180 degrees slew time to over 50 minutes on wheels around the largest inertia axis, and discourages the systematic use of propulsion for faster satellites slews.

Table 1: Main TGO satellite characteristics during aerobraking

Mass	1750 kg
Inertias	(4500, 5800, 2500) kg m ²
Solar Array total area	22.0 m ²
Total front area	29.3 m ²
Ballistic Coefficient	60 kg/m ²
Heat Flux Limit	2800 W/m ²
Heat Load Limit	500 kJ/m ²
Dynamic Pressure Limit	700 mPa
Reaction Control Thrusters configuration – all redundant	4x 10N pitch/yaw 4x 10N roll (balanced) 2x 10N force-free
Reaction Wheels configuration	4x 25 Nms (tetrahedral)
Worst-case slew time	54 minutes

Finally, the thermal and structural design impose survivability limits of 2800 W/m² heat flux, 500 kJ/m² heat load and 700 mPa dynamic pressure. These figures are considered hard limits and impose the choice of opportune safety margin in the design of the aerobraking corridor. The initial aerobraking mission design considered a 100% margin, leading to a nominal duration of less than 8 months, from Nov. 2016 to June 2017 [17]. After several iterations with project and industry and a recommendation by JPL, the margin was increased to 150% and the nominal duration increased to about one year from March 2017 to March 2018, including a two-month

interruption during a superior solar conjunction in July-August 2017. The first requirement on the navigation system consists on keeping the satellite within the predefined aerobraking corridor, ensuring the spacecraft safety for the entire duration of the prediction horizon, which is tuned depending on the aerobraking phase.

Although the correct TGO orientation during atmospheric passes is mainly achieved through passive aerodynamic stability, a thruster-controlled aerobraking mode is enabled, with a loose angular and rate dead-band that permits relatively large angular errors, up to 30 degrees in pitch and yaw axes and 15 degrees in roll axis, in order to limit the number of thruster actuations.

Some of the most important components of TGO autonomy are the Pericentre Time Estimator (PTE), based on accelerometer measurements, which estimates the time of the past two pericentres and predicts the upcoming one, therefore allowing the autonomous time-shift of the entire on-board commands timeline, including entry and exit from aerobraking mode; the autonomous protection manoeuvres, that reduce the aerobraking regime in case some of the aerodynamic constraints are close to the limits (Flux reduction manoeuvre, or FRM) or bring the satellite completely out of the atmosphere in case of a serious anomaly leading to safe mode (Pop-up Manoeuvre); a Pop-up Direction Estimator, that coarsely computes the direction of a pop-up manoeuvre based solely on accelerometer and gyroscopes measurements, in case all the star trackers fail to provide an attitude reading during the autonomous safe mode recovery; an Atmosphere Detector that enables a coarse attitude control if the atmosphere is entered with the wrong orientation during a safe mode. Several Failure, Detection, Isolation and Recovery (FDIR) functions are implemented to continuously monitor vital parameters and trigger immediate actions in case of anomaly. All these autonomous functions, although indispensable for the execution of aerobraking, needed complex trouble-shooting and fine tuning before being activated on-board [5].

One of the features that most affected the navigation of the satellite was the fact that all autonomous manoeuvres, planned to execute at apocentre, were based on a fixed table of times and directions, computed on-ground and reflecting the best prediction of the satellite ephemeris at the time of commanding. These predictions degrade quickly due to errors in the atmosphere density modelling. A large error between the predicted and actual phase of the spacecraft in its orbit would cause the manoeuvre to be less effective or even, in extreme cases, to be executed in a completely wrong direction. A protection against this eventuality was enabled, by monitoring on-board the difference between the pericentre time predicted by the PTE and the last prediction from ground, named DT_PTE. If a threshold of $1/8^{\text{th}}$ of the orbital period was violated, the FDIR would trigger a safe mode and execute a pop-up manoeuvre at the next apocentre. This sets a very stringent requirement on the navigation team, i.e. to ensure that within the scope of the next commanding window no violation of the DT_PTE parameter would occur. During the last months of aerobraking, this last requirement became the driver of the aerobraking corridor design, being stricter than the aerodynamic limits.

The inaccuracy of orbit predictions can also pose a threat to the spacecraft safety whenever its orbit crosses the one of another Martian satellite, natural or artificial. Orbit conjunctions need to be closely monitored and in case the relative distance reduces too much, a collision avoidance (COLA) warning is issued and formal actions must be taken, including the option to execute avoidance manoeuvres [18, 19].

The accuracy of on-ground orbit predictions is important also for the tracking of the satellite from ground antennas. Even though the pointing errors are small at Mars distance, the uplink frequency of the carrier signal needs to be continuously adjusted to pre-compensate the Doppler effect and guarantee on-board lock, which is needed for commanding and acquisition of two-way radiometric measurements. A large orbit phase error can prevent the on-board receiver to lock due to miss modelling of the Doppler shift. A robust strategy for acquisition needed to be developed, involving a large amplitude continuous frequency sweep profile to be superimposed to the uplink frequency profile. Details on this aspect can be found in Ref. 4.

No requirement existed on the accuracy of the orbit reconstruction, as no science activities were performed during the aerobraking phase. Even a loose position accuracy is considered sufficient to exploit the atmosphere density data gathered during aerobraking, which can be possibly used to improve the modelling of the Martian atmosphere.

Challenges on the Navigation Team

Interpreting the erratic and not well understood Martian atmosphere represents one of the most difficult challenges of aerobraking. The orbit-to-orbit variation of its density at the same height can easily exceed a factor two, and sudden, persistent onsets of up to 30% in the average value can occur at any time, quickly bringing the satellite to its aerodynamic limits. The navigation team must always guarantee the spacecraft safety, by constantly monitoring the status of the atmosphere and adjusting the aerobraking corridor to the new conditions, while also updating the relevant aerodynamic models.

Very frequent orbit re-planning and spacecraft commanding cycles need to be executed, including the design of pericentre correction manoeuvres. For the most relaxed phases, commanding is executed every two days, with an orbit prediction horizon of 3.5 days from the last tracking data used for orbit determination to the last command in the on-board nominal timeline (extra commands are available on-board to guarantee spacecraft safety in case of missed uplink). During the most critical phases, commanding cycles are run daily on extended working hours, and the prediction horizon is reduced to 2 days, further lowered to 1.5 days during end game (see Ref. 5 for details on Flight Dynamics commanding cycle). All operations like orbit determination, data analysis, decision making about model updates or control manoeuvres, and commanding are run under time pressure to meet the command delivery deadline, set to about 12 hours prior to the first command execution to allow for sufficient uplink time margins.

To increase the robustness of the navigation approach and reduce the risks, a collaboration program with the navigation team at the JPL was established, where they would share the expertise acquired during previous aerobraking campaigns and provide parallel Orbit Determination (OD) support during the most critical phases. OD solutions from the two teams were regularly exchanged and compared, and during the critical end game JPL provided daily solutions in-between ESA navigation cycles, acting in practice as a “night shift” with the possibility of contacting the ESOC team in case an emergency was detected.

The Orbit Determination Concept

The Orbit Determination of TGO is fully performed on-ground, relying primarily on 2-way Doppler observations acquired from the ESA and NASA deep space networks antennas.

During aerobraking, Doppler measurements were collected over at least two station passes per day, corresponding to approximately 16 hours/day, and increased to a full 24/7 coverage during the most critical last three months of operations, with orbit periods below 6h. The Doppler data are very sensitive to small errors of the satellite phase in its orbit, meaning that the ΔV s caused by drag can be accurately estimated, provided that a tracking pass is available once per orbit. For this reason Doppler data are very well suited for aerobraking, however if no tracking is available for three or more consecutive braking passes, the individual drag ΔV s can no longer be distinguished without recurring to other sources of information, e.g. accelerometer data collected on-board.

The dynamic models, especially those for the non-gravitational forces, were partly based on the a-priori information provided in the commanded sequences, and partly complemented with information obtained from the satellite telemetry, in a somehow hybrid setup. This was

important given that all commands were time-shifted by the PTE. In particular, telemetry information was used to identify the exact times of manoeuvre and wheel desaturation events, to model the ΔV s of unbalanced thruster firings controlling the attitude in aerobraking mode, and to access directly the accelerometers readings during atmospheric passes. The gravity model included point-mass contributions from the Sun, all the planets, Deimos and Phobos, plus the non-spherical harmonics expansion of Martian gravity from MRO120D model [20] up to degree and order 85, found to be a good compromise between accuracy and run time. The solar radiation pressure did not represent a major perturbation for TGO during aerobraking, thus it was modelled with a very simple flat plate model, with constant effective area and reflectivity coefficients, on top of which a linear scaling factor could be estimated by the filter. The atmospheric drag and the propulsion firings were the main non-gravitational perturbing forces and their modelling will be described in detail in dedicated paragraphs.

The satellite ephemeris is propagated starting from a specified epoch and covers the duration of the observation arc. A linear least square estimation filter, based on the Bierman's Square Root Information Filter formulation [21] and included in the ESA software package AMFIN [22, 23], is responsible of iteratively adjusting the parameters of the dynamic and observation models, up until the weighted root mean square of the measurements residuals (observed minus computed) is minimised.

The observation data span ends at the predefined data cut-off of the current planning cycle and starts shortly before the data cut-off of the previous one. This short overlap permits to compare successive solutions and generate a gapless orbit ephemeris, without increasing unnecessarily the length of the arc, thus reducing run time and filter convergence problems.

After the orbit is determined, the atmospheric density data collected during the recent past are post-processed and statistically analysed, in order to capture any atmospheric trends and update the atmospheric models before the re-definition of the aerobraking corridor in the next commanding period.

Modelling the Drag Force

During aerobraking, the atmospheric drag is the largest non-gravitational perturbation force acting on the satellite. The drag acceleration is computed with a very simple flat plate model:

$$\mathbf{a}_D = c_D p_{dyn} A \frac{\mathbf{V}_{in}}{\|\mathbf{V}_{in}\|} \quad (4)$$

Where c_D is the drag coefficient, p_{dyn} the dynamic pressure, A the exposed spacecraft area and \mathbf{V}_{in} the relative incoming velocity of the atmosphere flow. The complexity of the aerodynamic interactions between the atmospheric flow and the spacecraft is condensed here in the fixed parameter c_D , despite the fact that the loose attitude control dead-band permits errors of up to 30 degrees per axis to the nominal attitude profile. The on-board attitude, retrieved via telemetry, could have been used to create a more realistic aerodynamic model, however this would have increased noticeably the operational complexity, also creating a dependency that would impact the system robustness in case of telemetry outages.

Although this simplified model permits to capture the essence of the drag force and its effect on the orbit period reduction, it cannot fully represent the local deviations of the force magnitude and direction caused for example by atmospheric density variations, transverse aerodynamic forces, attitude errors or the presence of winds. A local force perpendicular to the incoming flow can arise: its component in the orbit plane will be called *lift*, the one perpendicular to it *side-slip*. Also a *time-shift* of the peak acceleration from the geometric pericentre can occur due to local atmospheric and topographic irregularities. The lateral force components can be large, reaching up to 10-20% of the main component, and can modify the

orbit orientation (inclination, line of apsides), for this reason they cannot be ignored in the orbit determination process. A very simple method, effectively employed during aerobraking, consists in modelling two impulsive ΔV s at the geometric pericentre, in the lift and side-slip direction, and a pair of equally sized and opposite ΔV s applied a fixed number of minutes before and after the pericentre, to align the satellite phase in the orbit, fixing the time-shift error. These impulses are treated as stochastic parameters and estimated by the orbit determination filter. The availability of the raw accelerometer measurements from the satellite telemetry also suggested a different and more direct route in the modelling of drag. The measurements could be pre-processed, in conjunction with the on-board attitude estimate, in order to remove outliers, biases and accelerations caused by attitude rates, and fed to the orbit determination in form of a reconstructed acceleration file in the inertial reference frame. More details on this approach, which was also implemented and validated in-flight, will be provided in a dedicated section. The difficulties related with drag acceleration modelling are exemplified in Fig. 1, where the variety of the acceleration profiles measured on-board is illustrated for a selection of braking passes.

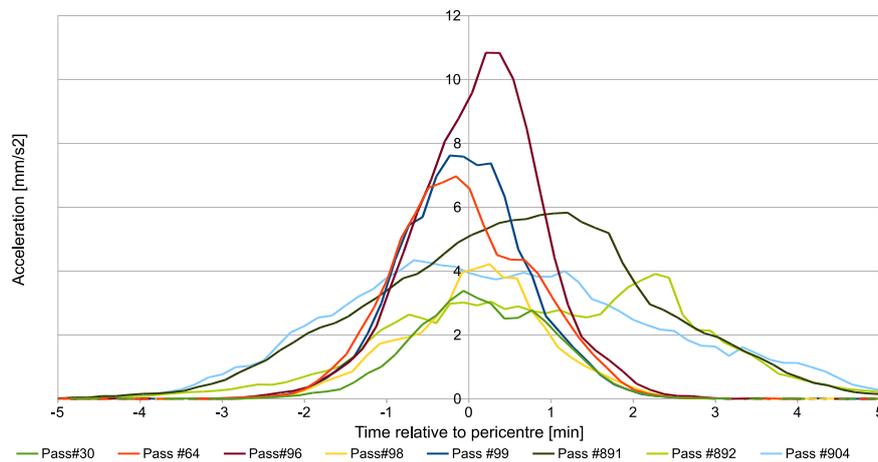


Fig. 1: On-board measured acceleration profiles for selected drag passes

Atmosphere Density Models

A reliable atmospheric density model is a key element to ensure accurate orbit predictions. The Martian atmosphere is far from being predictable, with orbit-to-orbit density variation that can easily exceed a factor two, as experienced during JPL aerobraking campaigns. Complex climatological models are now available, based on general circulation schemes and augmented by density data gathered by previous Mars missions; they should be capable of capturing seasonal and areographic density dependencies, even though local deviations are still very difficult to predict due to the lack of real-time data.

Three different models were implemented for TGO: the European Mars Climate Database (MCD) in its version 5.2, using the “climatology minimum solar conditions scenario” [24], NASA’s Mars Global Reference Atmospheric Model (MarsGRAM, 2010 version, [25]), configured as per JPL suggestions, and a simple exponential law, described in Eqn. 5:

$$\rho = \rho_0 e^{-\frac{h-h_0}{H}} \quad (5)$$

Where h is the height above the Mars reference ellipsoid, ρ_0 is the density at the reference height h_0 , and H is the so called scale height, a parameter describing the height difference causing a variation of a factor $e \cong 2.7$ in the density. The scale height is particularly important to size the corridor control manoeuvres, and needs to be monitored and updated whenever the

observed value deviates significantly from the reference model. Given two atmospheric passes with the same orbit and effective ΔV , the one with the lower scale height will present a higher peak acceleration and heat flux, thus it will bring the spacecraft closer to its aerodynamic limits. A linear density scale factor per pass is applied to any of the chosen models and estimated in the orbit determination in order to fit the observation data. A simple statistical analysis, based on a constant or linear regression, is run over a selected number of passes to monitor the trend of the density and correct the model for the future orbit prediction. The choice of the number of passes for the regression is a key element: too many passes result in slow and delayed model updates, failing to capture fast density variation trends and leading to biased predictions; too short regression spans can lead to oscillating updates, which only follow the noise of the density measurements. The optimum number of passes has been empirically found to be about 10-15, also depending on the orbit period and consequently the frequency of the atmospheric sampling. It must be remarked that a constant bias in the atmospheric density causes the phasing error to grow quadratically with the number of braking passes. Therefore, it is more important to predict correctly the average corrections to the density models than the local deviations.

Based on the evidence, gathered from previous JPL aerobraking mission, of persisting longitude-dependent density variations, a second and more complex way of making atmospheric prediction was implemented, based on a dedicated scaling function which included a linear function of time $SF(t)$ and a longitude-dependent cosine series $\gamma(lon)$ truncated to order three:

$$\rho = \rho_{model} \cdot SF(t) \cdot (1 + \gamma(lon)); \quad \text{with } \gamma = A_0 + \sum_{k=1}^3 A_k \cdot \cos(k \cdot lon + \varphi_k) \quad (6)$$

Where A_0 is a constant offset, A_k are the amplitude terms and φ_k the corresponding longitude phases. The amplitude and phase parameters could be estimated by fitting the recent history of atmospheric scaling factors. This function, informally called *wave model*, was tested starting approximately from the 6h orbit period, where at least 4 scaling factors per day could be collected. Although in some periods it demonstrated to improve the dispersion statistics of the predicted scaling factors, it was never used for the operational orbit prediction, as it did not bring significant benefit in terms of *accumulated* orbit phase error.

Modelling Thruster Actuations

Thruster actuations can be due to wheel desaturations, orbit control manoeuvres or attitude control. The nominal manoeuvres are planned ahead and can be time-shifted by the PTE, however autonomous manoeuvres can also happen.

A telemetry-based pulse file is derived from the on-board model and includes all attitude control thruster firings during aerobraking mode, expressed as a time series of impulsive ΔV s and already rotated to the inertial reference frame by making use of the on-board attitude estimate. The file is automatically pre-processed by a dedicated software and the information used in the OD process. The attitude control pulses are included in the OD without scaling, in the assumption that their modelling error is small and can be easily absorbed by the drag dynamic parameters, acting on the orbit at the same time. The parasitic ΔV s from wheel desaturations, executed in force-free configuration, are estimated as stochastic impulses at the time extracted from telemetry events. The ΔV s of orbit control manoeuvres are obtained from on-board accelerometer data, provided in a pulse file having the same format as that used for attitude firings, and scaled in three axes in the OD to account for accelerometer sensing and alignment errors.

In case telemetry is unavailable for any reason, control manoeuvres and wheel desaturations are modelled based on the nominal timeline and their exact timing is estimated by the OD filter,

while attitude control pulses during aerobraking mode cannot be modelled at all, which will cause errors in the estimate of the corresponding density scaling factors by up to 3-4%. This degradation is considered acceptable in a situation where the fluctuations of the density itself can be in the order of 100%.

Improving the Estimation Filter Convergence

The use of a linear estimation filter in a non-linear problem can cause filter instability issues, particularly due to the single shooting approach, where the satellite state vector is propagated starting from a single epoch. Especially toward the end game, every day several new orbit revolutions needed to be determined quickly, and the accumulated prediction errors prevented the filter convergence in a single OD run. A mitigation measure was identified in the iterative update of the a-priori values of the parameters to be estimated, in particular the drag scaling factors, after an intermediate solution with a shortened arc length. This process was partially automated and reliably guaranteed convergence in all cases with only 1-2 extra OD runs. Additional robustness was provided by the availability of the solutions generated overnight by the JPL navigation team, bridging more than half of the data gap from the previous navigation cycle: the a-priori values for the drag scaling factors and the spacecraft state vector could be possibly read from these solutions in case convergence difficulties were encountered.

Use of Accelerometer Measurements

TGO is equipped with two Inertial Measurement Units (IMUs) including linear accelerometers. Their measurements are used on-board to implement GNC autonomous functions, but are also accessible in the satellite telemetry. These measurements can be used in the estimation process in various forms to increase the accuracy and robustness of orbit reconstruction. The accumulated braking ΔV in the spacecraft or inertial reference frame can be directly obtained by integrating the measurements accelerations over a drag pass, once the measurements are processed in order to remove biases and correct for the effects of attitude rates and propulsion.

Pre-processing Techniques

Two different ways of including accelerometer data in the OD process were implemented, as described in the next paragraphs. Before their use however, a common pre-processing had to be applied to the raw data, recorded by the spacecraft at the 10Hz clock frequency and dumped to ground for the full period spent in aerobraking mode over each aerodynamic pass. The raw data were constituted by the acceleration vector measured by three orthogonal accelerometer channels, already rotated to spacecraft frame by the on-board software using the nominal alignment matrices, based on pre-launch measurements and with up to 0.5 degrees of possible direction errors.

The output of the pre-processing was an acceleration pulses file, containing ΔV impulses in inertial frame, cumulated from the acceleration values over a configurable period, selected to be of 10 seconds as a compromise between filtering and resolution requirements.

As a first pre-processing step, the accelerometer biases in each spacecraft axis were removed, by subtracting from all 10Hz measurements the average of the acceleration readings in 10-min intervals before and after the aerodynamic pass. Linear fit of the pre- and post-pass biases were also attempted, but did not result in sufficient improvements to justify the lower robustness to spurious acceleration readings outside of the atmosphere (e.g. thruster actuations). As a second step, the acceleration vectors were converted for each data point from spacecraft to inertial frame, using the attitude quaternion from telemetry as measured by the on-board Gyro-Stellar

Estimator. The resulting acceleration vectors were then summed over consecutive 10-second intervals $t_1 \rightarrow t_2$, resulting in the ΔV s in Inertial frame sensed at the IMU location for each cumulated period, $\Delta \mathbf{V}_{IN}^{IMU}$. As last step, the effect of angular rates was removed to obtain the ΔV s in Inertial frame at Centre of Mass (CoM) $\Delta \mathbf{V}_{IN}^{CoM}$ from the $\Delta \mathbf{V}_{IN}^{IMU}$, as follows:

$$\Delta \mathbf{V}_{IN}^{CoM}(t_1 \rightarrow t_2) = \Delta \mathbf{V}_{IN}^{IMU}(t_1 \rightarrow t_2) - \left(\mathbf{q}_{SC \rightarrow IN}(t_2) \circ (\boldsymbol{\omega}_{SC}(t_2) \times \mathbf{P}_{SC}^{CoM \rightarrow IMU}) - \mathbf{q}_{SC \rightarrow IN}(t_1) \circ (\boldsymbol{\omega}_{SC}(t_1) \times \mathbf{P}_{SC}^{CoM \rightarrow IMU}) \right) \quad (7)$$

Where $\mathbf{q}_{SC \rightarrow IN}$ is the quaternion from Spacecraft to Inertial frame, $\boldsymbol{\omega}_{SC}$ are the angular rates in Spacecraft frame as measured by the gyroscopes, and $\mathbf{P}_{SC}^{CoM \rightarrow IMU}$ is the CoM to IMU vector in Spacecraft frame. Note that this is equivalent to subtracting the instantaneous accelerations due to rotational motion in Spacecraft frame before rotation and integration. The resulting ΔV s at the CoM were finally written to the output pulse file used by the OD, together with the related intervals start and mid time. Note that cumulating over the given intervals leads to filtering both the accelerometer noise and several high frequency phenomena related to atmospheric fluctuations, spacecraft appendages flexible modes and thruster actuations. The integrated values however correctly represent the total ΔV s, with no significant effect on the OD process introduced by the loss.

As an example of the pre-processing, Fig. 2 shows the measured acceleration along the X axis of the spacecraft for pass #161 on the 19th September 2017, which resulted in the largest ΔV of the whole aerobraking, triggering a Flux Reduction Manoeuvre. The processed cumulated ΔV pulses are converted back to constant accelerations in Spacecraft frame shown in black, and compared to the raw data and to a smoothed profile obtained by applying a 10 s Gaussian filter.

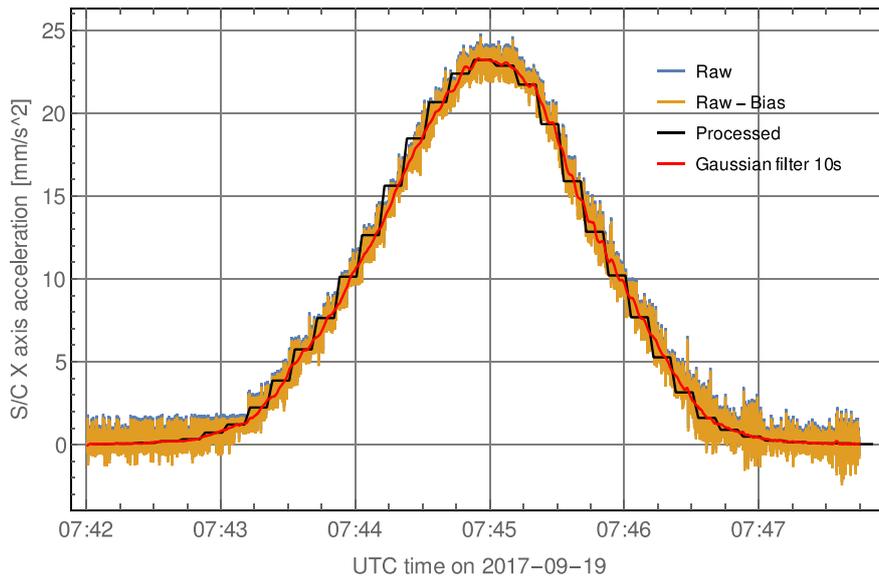


Fig. 2: Example of raw and processed accelerometer data from atmospheric pass #161.

Note that telemetry monitoring and processing activities on Flight Dynamics extended much beyond the generation of the OD input files described in this paper (spacecraft events, thruster ΔV pulses, accelerometer ΔV pulses). A broader description of these activities, among which the assessment of aerobraking propellant consumption, attitude control bands and aerodynamic stability, the independent measurement of ΔV s to be compared with the OD results, atmosphere scale height pass-per-pass monitoring, etc., are described in more details in Ref. 5.

Accelerometer Measurements as Dynamic Models

A possible way to include the measurements in the orbit determination is to transform them into an acceleration profile, already rotated to the inertial reference frame, and feed them directly into the orbit propagator as a dynamic model. This method has been implemented in the OD software and was used in-flight as the baseline drag model during the first months of aerobraking.

The main advantage of this approach is that it removes almost completely the need to model complex aerodynamic interactions during orbit reconstruction. The acceleration profile already includes the effects of attitude oscillations during braking mode, thus correctly modelling lift, side-slip and peak acceleration time offsets. A single linear scaling factor per pass was still estimated to correct for possible measurement and pre-processing errors, resulting in corrections that did not exceed 0.5% at nominal aerobraking altitude.

On the downside, the reconstructed orbit needed an extra post-processing step to convert the acceleration or the integral ΔV information into a scaling factor over the reference atmospheric model, to be used for statistical analysis and consequently for the orbit predictions. Additionally, some evidence of accelerometer misalignment was noticed during higher aerobraking regimes, causing direction errors in the acceleration model and some stiffness in the measurements fit, visible from the observation residuals. This kind of error was difficult to correct directly in the OD process with the available software, and would have required additional software developments.

Accelerometer Measurements as Observations

A second way to include accelerometer measurements in the OD is to use them as observations in the navigation filter. Although all aerobraking satellites so far were equipped with accelerometers, the direct use of these observations in the orbit determination presented practical difficulties, in particular serious convergence issues caused by the very narrow linearity region of the observation models, as detailed in [26]. A different approach, consisting in the use of acceleration measurements integrated over an entire atmospheric pass, was under development at JPL and was also implemented by the ESOC team.

This type of observation is much less sensitive to the accumulated orbit prediction errors and can complement Doppler data, for example allowing to distinguish the individual passes ΔV s in the case of a prolonged gap in the radiometric tracking data.

If N accelerometer measurements are collected over a pass, being dt the sampling frequency and $d\mathbf{V}_i$ the integrated acceleration over one sample, the observed $\Delta\mathbf{V}_o$ vector can be obtained simply by summing all the measurements:

$$\Delta\mathbf{V}_o = \sum_{i=1}^N d\mathbf{V}_i \quad (8)$$

An equivalent vector $\Delta\mathbf{V}_c$ must be computed in the OD software based on the propagated trajectory and the acceleration models \mathbf{a} of the atmospheric drag and other perturbations:

$$\Delta\mathbf{V}_c = \int_{t_0}^{t_f} \mathbf{a} dt \quad (9)$$

Where the times t_0 and t_f are fixed times, defined at the time of measurement pre-processing by applying a fixed offset before and after the centroid time of the acceleration profile.

The software permits to model all non-gravitational forces, including propulsion, however for TGO its usage was restricted to the sole drag model; propulsion accelerations, computed from the on-board model, were subtracted from the accelerometer measurements during pre-processing. To include the observation in the OD filter, the drag acceleration and its partial derivatives with respect to all dynamic parameters needed to be computed during orbit propagation and saved to a temporary file; eventually, the integrated quantities were computed by quadrature. A simple trapezoid method with fixed time step proved accurate enough to avoid numerical issues. Any impulsive ΔV , like the ones used to model lift, sideslip and time shift, was later summed to the observation model, together with its partial derivatives.

A preliminary version of the observable was initially implemented by taking the norm of the $d\mathbf{V}_i$ vectors *before* summing. This however presented two major disadvantages: even after pre-processing and de-biasing, the residual noise of the measurements accumulates when summing vector magnitudes, deviating from the hypothesis of Gaussian distribution of the measurement errors; furthermore, it was difficult to find a correct mathematical formulation to include discrete impulses in the observation model. Retrospectively, both issues could have been sorted by modifying the observation definition, and applying the norm operator *after* the integration, as shown in [26].

The vector definition does not present this problem, however it generates three measurements that are not statistically independent to each other, especially when rotated to the inertial frame. When weighting the measurements in the filter, this correlation was ignored for simplicity, accepting some inaccuracy in the computation of the a-posteriori orbit uncertainties.

The new software implementation was initially validated through cross-verification tests by comparing OD solutions with the JPL team, and later operationally validated by using real aerobraking data: a normal routine solution was compared with a case where a long Doppler outage was simulated, and the accelerometer measurements were used to fill the gap.

Orbit Predictions and Corridor Control Design

The initial design of the reference trajectory, called Long Term Plan (LTP) was based on the maximum possible ΔV that satisfied the aerodynamic constraints. The pericentre altitude profile was optimised to keep a 150% margin on the most demanding constraint, being heat flux for higher orbit periods and heat load for lower periods. The safety margin was initially applied on the *maximum* heat flux or heat load based on the MCD density model. This model however attempts to capture large longitude dependent density variations, making the LTP design somehow too conservative and the aerobraking progression relatively slow (the margin on the *average* dynamic constraints resulted to be more than 150%).

At short term planning level it was realised that the density variations predicted by the MCD did not find confirmation in the actual drag experienced by the satellite. This resulted in a very difficult interpretation of the density scaling factors estimated in the orbit determination. Shortly after the end of the walk-in, at the beginning of May 2017, the reference model was changed to the simpler exponential model, easier to interpret and update. At the same time the aerobraking regime was increased, by setting the *average* heat flux to $\sim 1200 \text{ W/m}^2$, thus slightly higher than nominal target value, to recover the accumulated delays caused by a colder than expected atmosphere and the troubleshooting of some spacecraft autonomy weaknesses.

The aerobraking progression, in terms of flown apocentre height against the one predicted at LTP, is illustrated in Fig. 3 for the entire duration of the aerobraking. The flat region in July-August corresponds to a superior solar conjunction, in which aerobraking was interrupted and the LTP trajectory was redefined, as visible from the jump between the blue and orange curves.

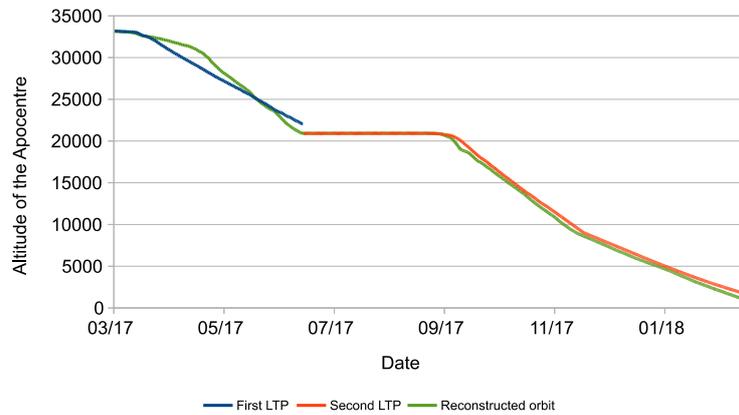


Fig. 3: Evolution of pericentre altitude: long term prediction vs. reconstruction.

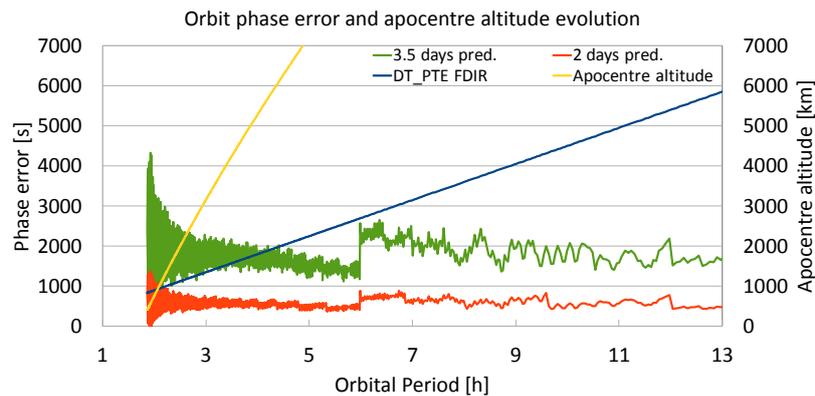


Fig. 4: Analysis of on-ground phase prediction errors, assuming 30% constant error in the atmospheric density.

At this stage, it was analysed that the DT_PTE would be the constraining parameter of the corridor design for the last stages of aerobraking. In Fig. 4 the predicted phase errors with a prediction horizon of 2 days (daily commanding) or 3.5 days (commanding every second day) are reported in the worst-case assumption of a persisting density error of 30%. The blue line represents the FDIR limit for the DT_PTE parameter, equivalent to $1/8^{\text{th}}$ of the orbit period. To prevent the possibility of a safe mode triggering, the aerobraking regime was reduced by 30% when reaching the 6h orbit period, while daily operations started shortly before 5 hours period was reached. It was also decided that aerobraking would terminate when the prediction error would reach the FDIR limit for a 2-day prediction horizon with a persistent 30% prediction error, corresponding to an altitude of 1000 km. The science orbit acquisition would then be completed with chemical manoeuvres, permitted by the fuel budget margin.

Close to the end game, further measures were implemented to improve the robustness against the possibility of a DT_PTE induced safe mode. The most effective was a so called “mini commanding cycle” executed after a new orbit prediction was made available. The apocentre tables for the autonomous manoeuvres were updated on-board already 12 hours earlier than the full sequence of commands, and the thresholds on the DT_PTE asymmetrically tuned accordingly, to correct the already known part of the prediction error. This shortened the prediction horizon from 2 days to approximately 1.5 days, which is a significant improvement given the quadratic growth of the errors with the number of passes.

During short term prediction cycles, the decision on the execution of corridor control manoeuvres was taken based on the most recent atmospheric density observation. Heuristic criteria more than a rigorous rule were used, such as following the average ΔV per pass from the LTP, avoid manoeuvres resulting in altitude changes less than 0.5 km and delay or advance manoeuvres in presence of COLA warnings.

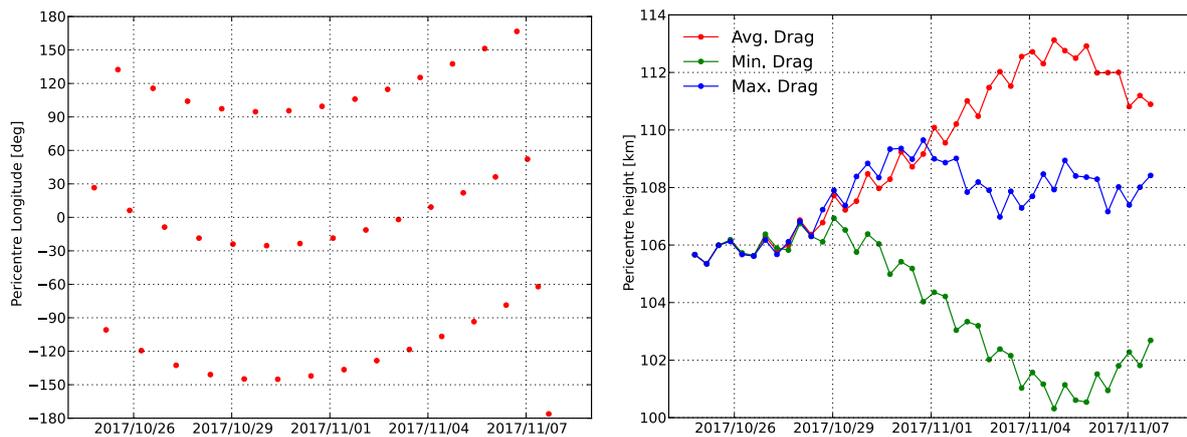


Fig. 5: Pericentre longitude evolution during 3:1 resonance with Mars rotation (left); Different evolution of pericentre height for different atmospheric density scalings (right).

The design of these manoeuvre was complicated by the interaction between gravitational and atmospheric effects on the orbit. Gravitational kicks can change the altitude of pericentre by more than a kilometre in a single pass, and this can be amplified by resonances between the orbit period and the Mars rotation period, when the same longitudes are flown several times. The exact longitude flown depends on the phase in the orbit, which is instead mainly dictated by the drag and cannot be known accurately in advance. An example of such a scenario is given in Fig. 5 for the 3:1 resonance in October 2017: three different scalings for the predicted atmosphere (nominal and worst case maximum and minimum) could lead to very different pericentre height evolution if the orbit is not controlled. In such a situation, daily pericentre control manoeuvres were necessary to maintain a constant aerobraking regime.

The risk of collisions with other satellites, both natural and artificial, was monitored independently at JPL, using the already existing MADCAP monitoring system [18], and at ESOC by the means of dedicated software tools. The long term evolution of TGO orbit was mainly driven by gravity perturbations, thus the close approach “seasons” were approximately known already at LTP level, provided that TGO would remain close to the reference trajectory. The orbit conjunctions close to apocentre were difficult to treat, due to the large uncertainty on apocentre height and phase, and the impossibility to foresee all scenarios, including the possible triggering of autonomous manoeuvres at any orbit. In particular, the case of Phobos conjunction in November 2017 required special care, considering also the large size of the Martian moon, and the maximisation of the orbit phase separation between Phobos and TGO was enforced in the fine tuning of corridor control manoeuvres over the last week before conjunction. The conjunctions close to pericentre were easier to treat, since the pericentre height is easier to predict and less sensitive to density variations. In this case, a safe orbit separation can be enforced with a relatively inexpensive collision avoidance manoeuvre aimed at increasing the radial distance. This strategy was used for example during the first walk-in, when a PLM was advanced by two days to avoid a close approach with Mars Express on 16th April 2017.

In-flight Results

The TGO aerobraking campaign lasted almost one full year, from 15th March 2017 to 20th February 2018. Operations were interrupted for two months, from 26th June to 30th August 2017 during a superior solar conjunction. A short summary of the main aerobraking achievements is given in Table 2. The aerobraking proved to be overall fairly efficient, with a net drag ΔV about 17 times larger than the one paid by the propulsion system.

Table 2: TGO aerobraking summary

Total duration including conjunction	342 days
Number of drag passes	952
Orbit period reduction	from 24 h to 2.1 h
Total ΔV from drag	1017 m/s
Total ΔV from attitude control	9 m/s
Total ΔV for manoeuvres	51 m/s
Number of PRM including walk-out	22
Number of PLM including walk-in	46
Number of autonomous FRM	1

Fig. 6 presents the evolution throughout aerobraking of several quantities of interest. In plot a) the evolution of pericentre altitude is depicted, together with all pericentre control manoeuvres including walk-in and walk-out. It is interesting to note how the natural pericentre height oscillations caused by gravity perturbations became wider with the progression of aerobraking, with orbit-to-orbit variations that could be larger than a kilometre. The dense accumulation of PLMs at the end of October 2017 was necessary to keep a roughly constant height during the 3:1 resonance between the TGO orbit and the Mars rotation periods. During mid-November, when reaching the 6h orbit period, the aerobraking regime was gradually reduced by raising the pericentre by about 3 km, in order to improve prediction errors. A constantly decreasing trend in the density toward southern latitudes required to progressively reduce the height again in December-January. Unexpectedly, the density did not raise with the same rate passing again through the same latitudes northbound.

In plot b) of Fig. 6 the history of the density scale factors reconstructed in the orbit determination is shown, together with the corresponding predicted scaling. It can be seen that during the first walk-in the prediction capabilities were relatively poor, due to a combination of factors. The MCD model used at the time was predicting large density variations that did not find correspondence in the reconstructed densities. Moreover, the unavailability of a sufficient number of samples at each altitude made the statistical analysis too much dependent on data acquired in the past and at different heights, which caused delays in capturing the decreasing trend of the density. These difficulties suggested the adoption of a simpler model, a reference exponential model with scale height 8 km, which at least allowed a better understanding of the density trends.

During the second walk-in, around mid-September 2017, the scale height of the reference exponential model proved to be no longer adequate to correctly control the aerobraking corridor. A relatively fast walk-in strategy had been designed, which did not permit to observe more than 1-2 atmospheric passes before computing the next PLM. After the last PLM was commanded, it became clear that the actual density increase was higher than predicted by the model and the heat flux was approaching the FDIR limit, at the time set to 2240 W/m^2 , which represents a 25% margin on the satellite limits. This can be seen in plot c) of Fig. 6, which reports the history of the peak heat flux. An autonomous Flux Reduction Manoeuvre was triggered after pass #161, immediately followed by an already commanded PRM. The heat flux slightly exceeded the specification, at about 2900 W/m^2 , although without any consequence on the spacecraft. A retrospective analysis on the previous passes, including the ones before solar conjunction, showed that some indications of a lower scale height were already present. In particular, the peak acceleration measured by the accelerometers was systematically higher than the one reconstructed based on the scaled model, as illustrated in plot d) of Fig. 6. A trend of the ratio between measured and modelled acceleration started immediately after the switch to the exponential model and settled after June 2017 to an average value of about 1.2. By fitting a scale height parameter for all the previous individual aerobraking passes, a value of 6 km was found to be more representative of the current conditions and the reference exponential model was updated on 18th September 2017.

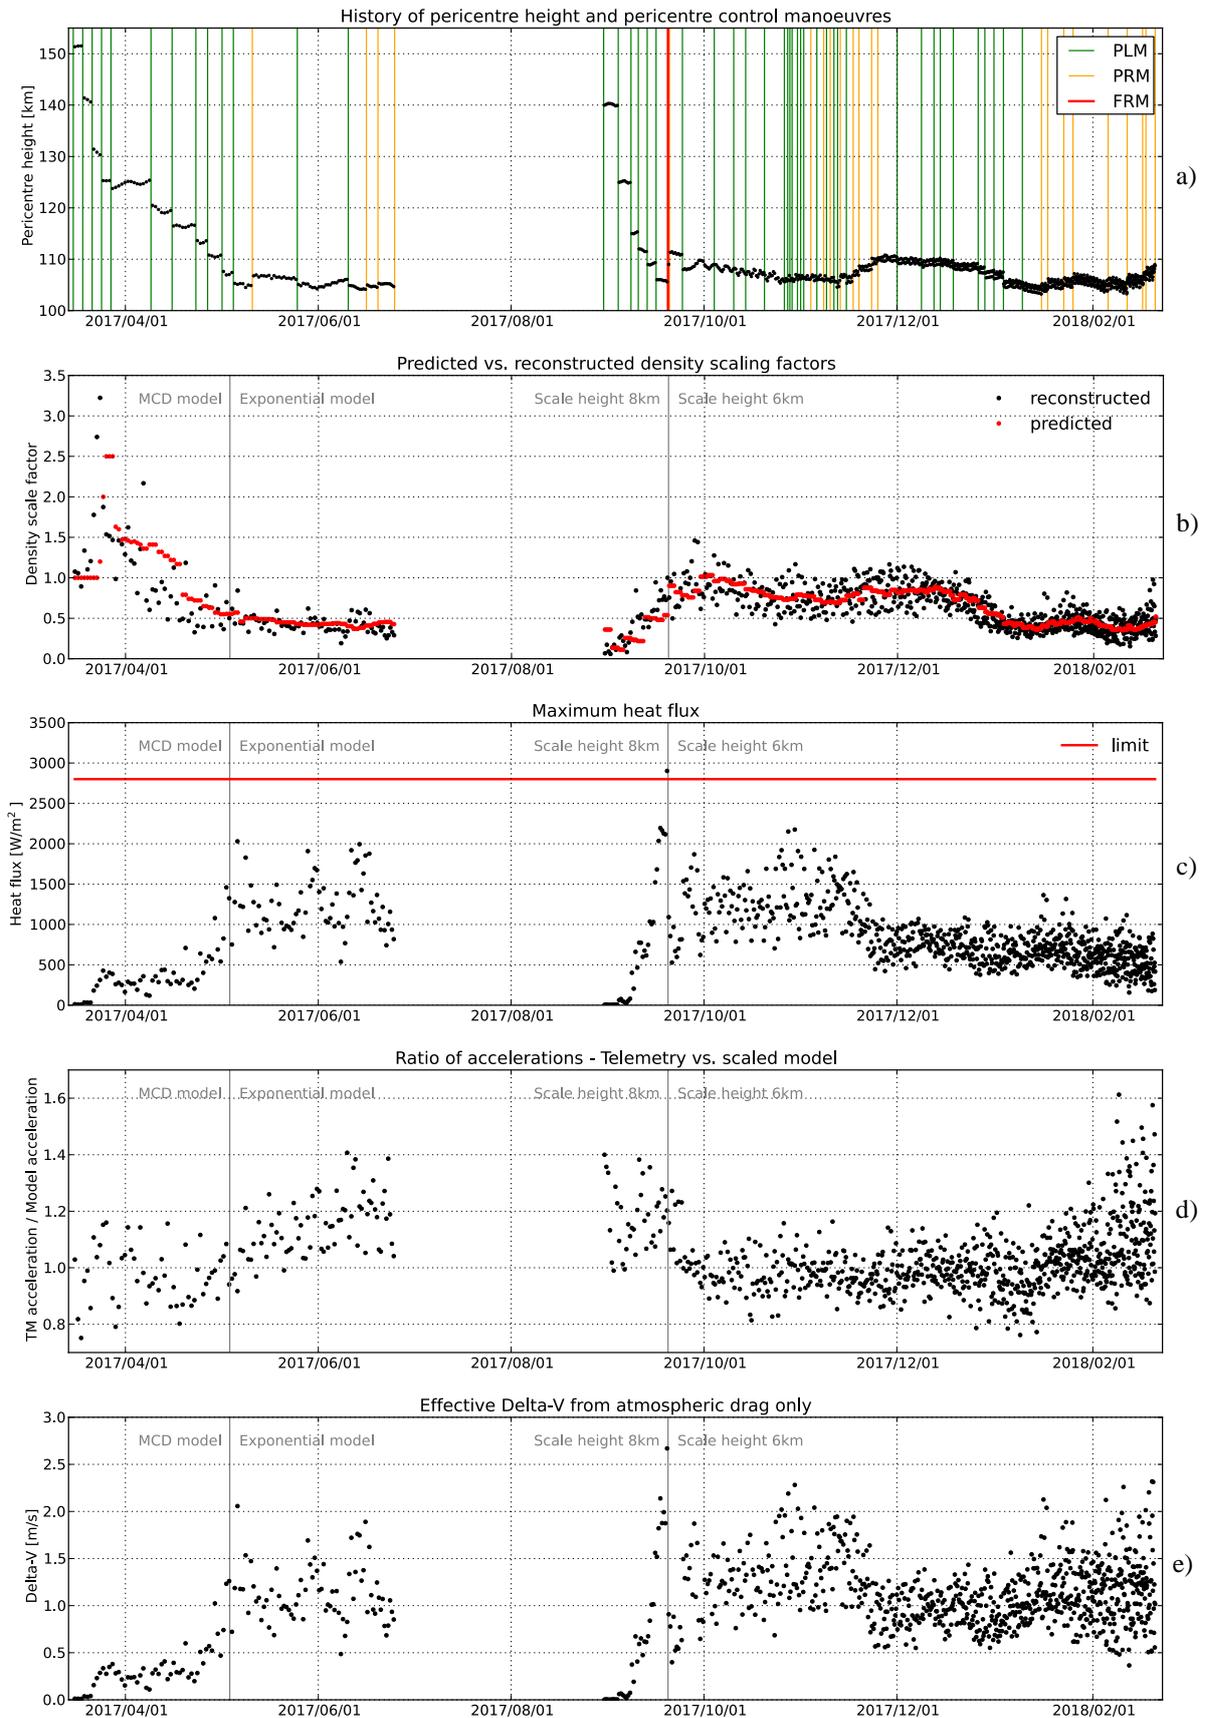


Fig. 6: Summary of aerobraking quantities of interest: a) Pericentre height and control manoeuvres; b) Predicted and reconstructed atmospheric density scale factors; c) Peak heat flux based on scaled atmospheric density model; d) Ratio between peak acceleration from telemetry and from scaled atmospheric model; e) ΔV per pass, due to drag only.

A strict monitoring of the scale height was implemented since, however the new value remained adequate up until February 2018, when the scale height further decreased. At that time the peak heat flux and dynamic pressure were no longer a concern, given the reduced aerobraking regime and the longer, shallower atmospheric passes, and the exponential model was not updated any further.

Finally, plot e) of Fig. 6 reports the history of the ΔV produced by the sole atmospheric drag, with the average value at regime being between 1-1.5 m/s, and a peak value of 2.7 m/s.

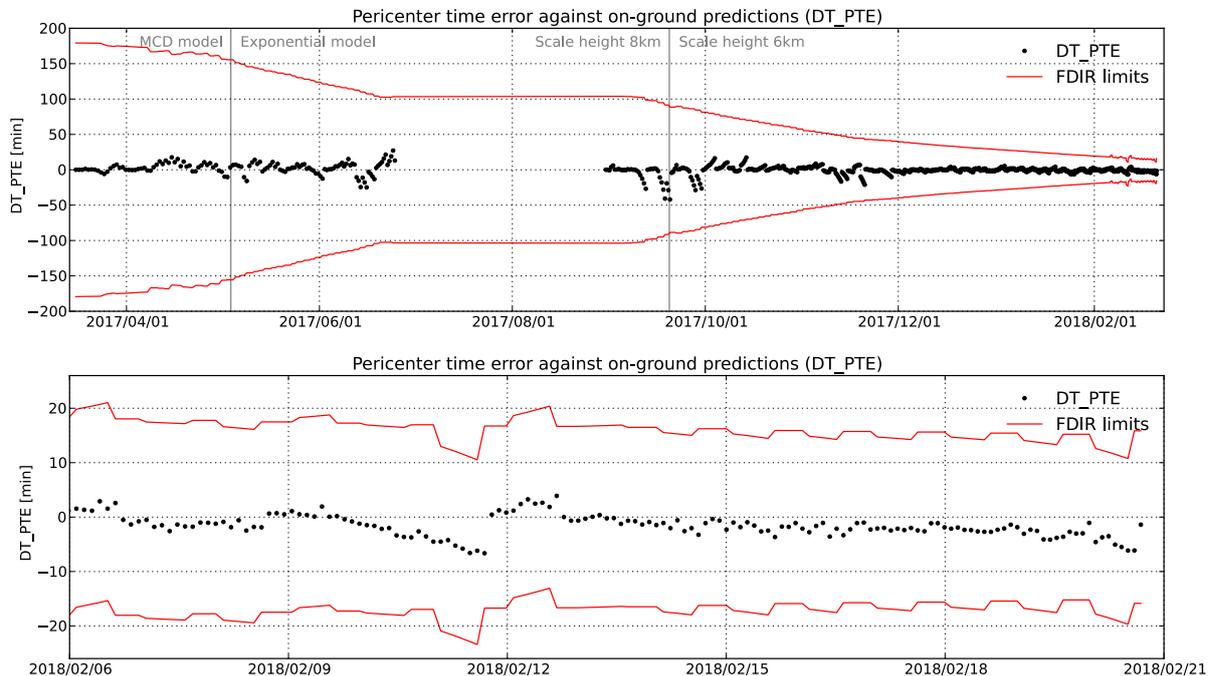


Fig. 7: Evolution of DT_PTE, bounded by the FDIR limits. Bottom plot: detailed view over the last two weeks, showing asymmetric bounds and improved predictions due to mini-cycle.

As explained in the introduction, the accuracy of the orbit phase predictions had a vital importance and could be directly monitored on-board and on-ground through the parameter DT_PTE, i.e. the difference between ground-predicted and PTE predicted pericentre time. The evolution of DT_PTE is illustrated in Fig. 7, together with the FDIR limits, set to $\pm 1/8^{\text{th}}$ of the orbit period. Although the triggering of a DT_PTE safe mode was never too close, in a few occasions the error reached about half of the threshold. The largest absolute error of about -42 min occurred right after the triggering of the autonomous FRM in September 2017.

In the zoom view of Fig. 7 the effectiveness of the previously described commanding mini-cycle is shown. About 12 hours before every mission timeline update, the tables containing the autonomous manoeuvre times were updated on-board, together with asymmetric shifts of the FDIR DT_PTE limits, accounting for the already-known portion of prediction errors. This measure guaranteed good margins during the entire duration of the end game.

Orbit Determination Results

The OD system behaved very reliably during the entire duration of aerobraking. Even though the setup was designed to be robust against partial telemetry or tracking data outages, the telemetry was always made available and pre-processed before the OD was run. When the orbit period reduced and the number of new orbits to be determined in a single solution increased, some convergence difficulties appeared in the estimation filter due to the non-linearity of the problem. The iterative approach described earlier however proved very effective, allowing in all cases to achieve convergence quickly.

Accelerometer observations were continuously monitored during the second half of aerobraking, by checking their residuals against an orbit determined with Doppler only. Orbit solutions were periodically generated offline including also the accelerometer observations and compared to the Doppler only ones, showing very good consistency. An example of Doppler and accelerometer measurements residuals for such an OD is shown in Fig. 8. The accelerometer residuals are in line with the adopted weighting of 5 mm/s, not too aggressive to accommodate for possible errors due to uncorrected biases and misalignments. The orbit comparison with the operational, Doppler-only solution, illustrated in Fig. 9, shows that the differences are mainly in the cross-track direction, the one worst resolved by the Doppler observations, while the in-track difference remains within a few meters, meaning that the Doppler data provide already very good accuracy in the orbit period reconstruction.

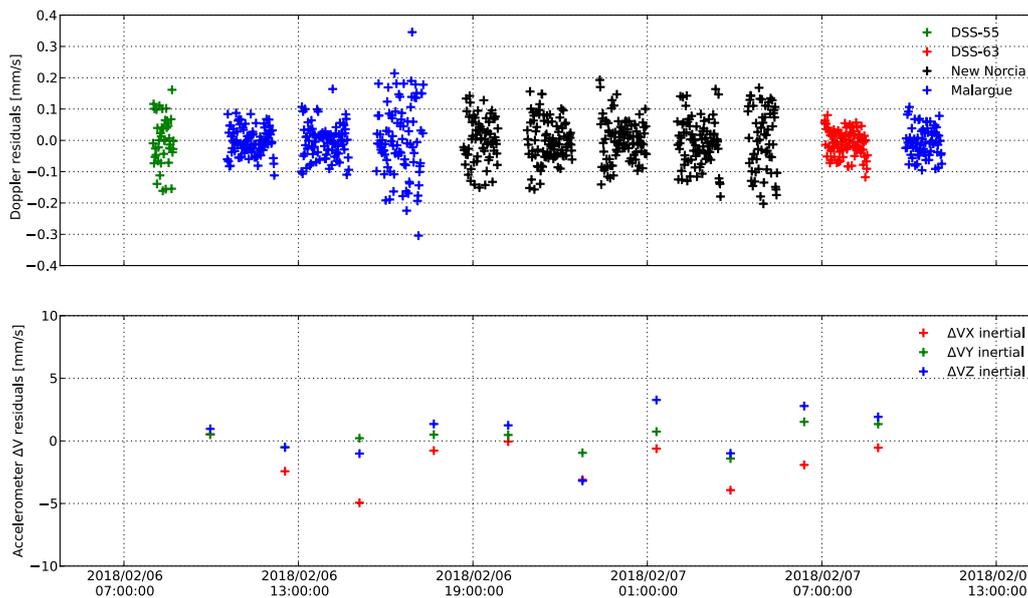


Fig. 8: Example of OD post-fit observation residuals, off-line reconstruction based on Doppler and accelerometer measurements.

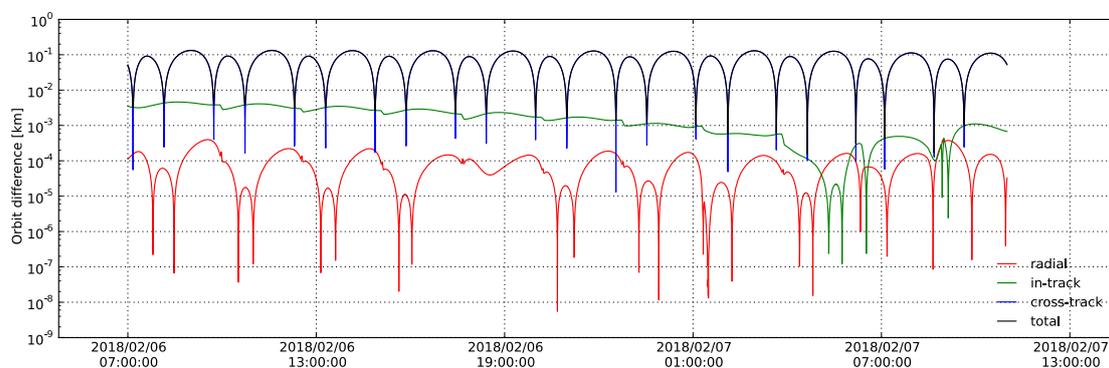


Fig. 9: Position difference between operational Doppler-only and offline orbit reconstructions.

In a single situation during operations, a tracking data outage at New Norcia station lasted four consecutive drag passes, and the operational use of accelerometer observations was made necessary to recover the individual atmospheric scaling factors and provide a good orbit reconstruction. The OD residuals in Fig. 10 show how the accelerometer data, which were downlinked during the following pass at Malargue station, filled the Doppler data gap that lasted about 11 hours.

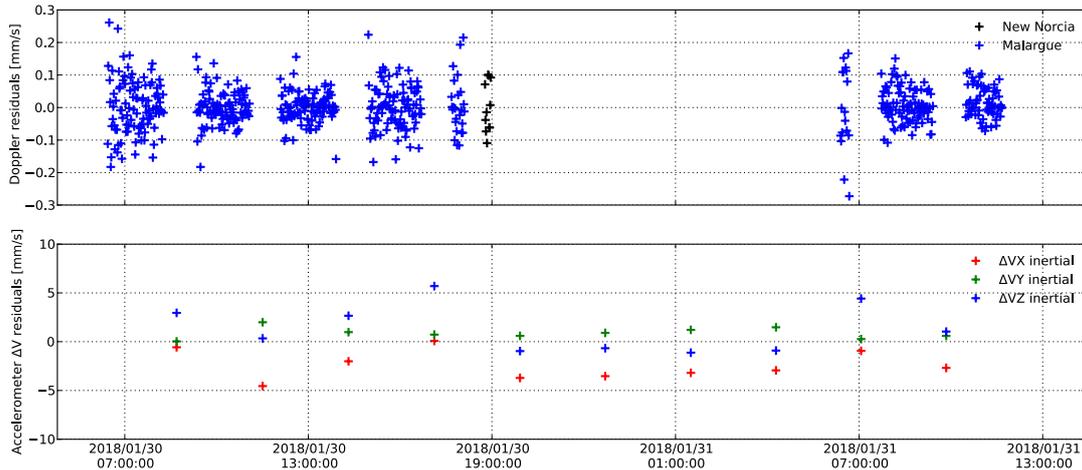


Fig. 10: Post fit observation residuals for an operational OD using accelerometer observations in presence of a station outage.

Atmospheric Models Performance

TGO provided a large amount of data by sampling the upper layers of the Martian atmosphere with high accuracy accelerometers in a wide range of latitudes, between -10° and -74° . These data will be used by the science community to further improve the current climatological models, but are also useful to evaluate the effectiveness of different models in spacecraft operations context, where the primary requirements are ease of use and interpretation and good orbit prediction capabilities.

A useful statistical indicator for predictability, called *persistence* p^N and defined as the ratio between the current scaling factor and the average of the previous N , was used during aerobraking to evaluate atmospheric models. When the number of previous passes is only one, the persistence coincides with the pass-to-pass variability p^1 , displayed against pericentre latitude throughout aerobraking in the left plot of Fig. 11, together with its 30-pass running standard deviation. The black curves and points are south-bound and the red north-bound.

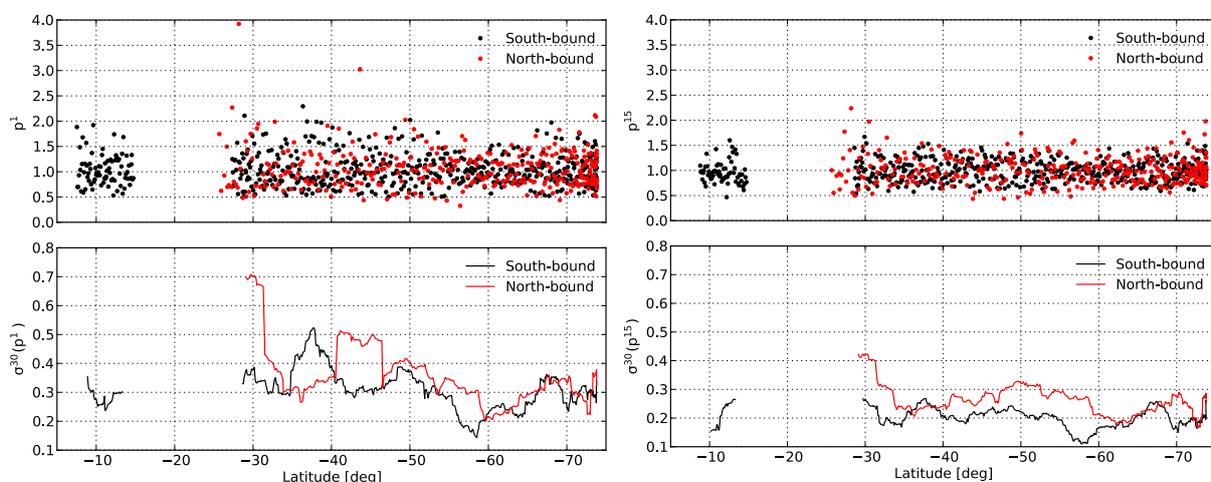


Fig. 11: Pass to pass variability p^1 (left plot) and 15-pass persistence p^{15} (right plot) against latitude of the density scale factors for the exponential model, together with their running standard deviations computed over 30 passes.

It can be seen that frequently the p^1 exceeded a factor 2 (or equivalently was less than 0.5), and in one extreme case even reached a factor 4. The overall standard deviation of p^1 is 0.35,

consistent with those measured by Mars Global Surveyor (0.39), Odyssey (0.47) and Mars Reconnaissance Orbiter (0.36) [27]. The persistence p^{15} computed over 15 passes and shown in the right plot of Fig. 11 gives a better perspective of how the predictability can change at different latitudes, or even at the same latitudes but at different epochs and conditions, as clearly visible from the separation of the red and black standard deviation curves.

As an extension of the concept of persistence, a *prediction function* can be defined. If the mean μ_i of the scaling factors SF_i over the last N passes is used for the prediction (Eqn. 10), the prediction function pf_i is the root mean square of the scale factor error of the following M passes, normalised by μ_i (Eqn. 11).

$$\mu_i = \frac{\sum_{j=0}^{N-1} SF_{i-j}}{N} \quad (10)$$

$$pf_i = \frac{1}{\mu_i} \sqrt{\frac{\sum_{j=1}^M (SF_{i+j} - \mu_i)^2}{M}} \quad (11)$$

This function is chosen to compare the predictive capability of different atmospheric models because it mimics the operational procedure used to update the predicted density scale factor, first running a regression over a number of past passes and then propagating the obtained scale factor to the future, and because it can combine in a single metric the effect of density trends (e.g. a continuous decrease, which would result in a bias) and of the pass-to-pass scatter.

Fig. 12 shows the evolution of pf_i for the exponential, MarsGRAM 2010 and MCD models during the south-bound leg of the aerobraking, with $N=M=15$. It can be seen that the prediction function for all models follows a similar trend, dictated by the local and short time scale density variations that are not captured by any of the models. During limited periods of time at mid latitudes the MCD seems to provide slightly better statistics, visible from the green curve below unity in the bottom plot, normalised to the exponential model. This is possibly due to its internal representation of stationary density waves getting periodically in phase with the actual ones; however the MCD performance deteriorates when entering the polar region, south of -60 degrees. The MarsGRAM prediction function oscillates around the one of the exponential model during the entire period, with higher amplitude at the southernmost latitudes.

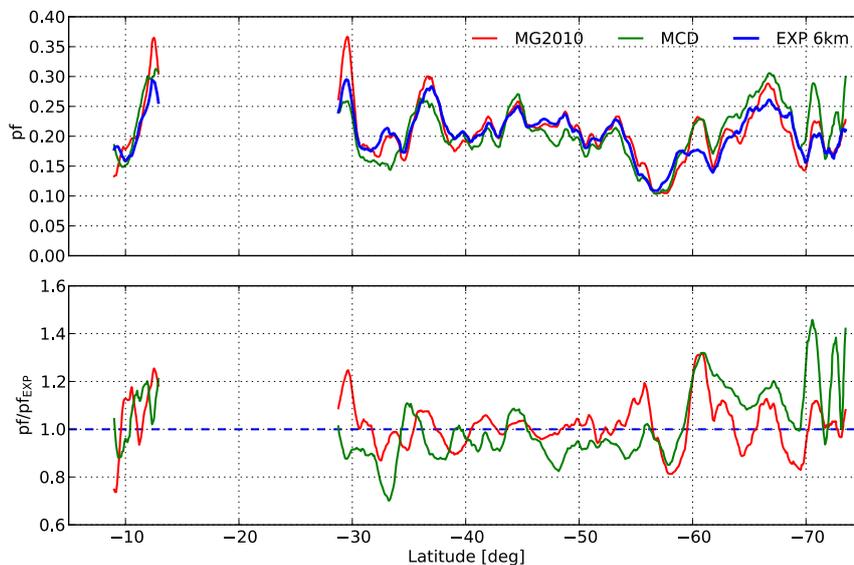


Fig. 12: Comparison of the prediction function for different atmospheric models along the south-bound part of aerobraking. Bottom plot: normalised to the exponential model.

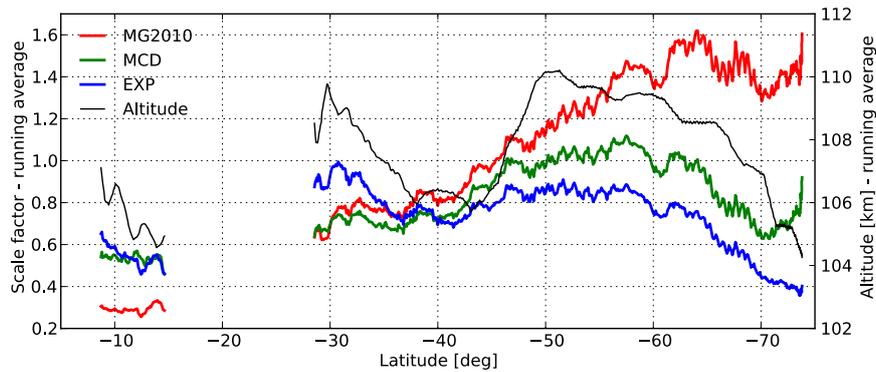


Fig. 13: Evolution of the atmospheric scaling factor vs. latitude for different models along the south-bound part of aerobraking, together with pericentre altitude plot.

In Fig. 13 the running average over 15 passes of the density scale factor vs. latitude is presented for the three models, together with the pericentre altitude, for the south-bound leg. The difference in the slopes of the different curves exemplifies how differently the models predict latitude-dependent density trends: MCD and especially MarsGRAM predict a significant decrease in density between 35 and 60 degrees latitude, which is not observed in reality. Below -60 degrees the polar region is entered, with a significant drop in atmospheric density that both MCD and MarsGRAM predict to happen at a lower latitude, around -70 degrees.

These metrics, although not fully representative of the complexity of the problem, give an idea of how the more sophisticated circulation models can qualitatively provide good descriptions of general atmospheric phenomena, but at the same time may fail to provide better quantitative predictions in the short term when used for spacecraft operations. The fact that no density model proved to systematically outperform the others during the entire TGO aerobraking led to the conclusion that the exponential model was operationally preferable due to its inherent simplicity.

Conclusions and Lessons Learnt

TGO was the first ESA satellite to successfully complete an aerobraking campaign. The extensive preparations from all the involved parties, a robust and carefully validated operational approach and a very large working effort of the mission control team, together with a very good performance of the satellite itself, permitted to run aerobraking operations very smoothly and to reach the target orbit slightly ahead of schedule, without any safe mode or serious contingency. From the Flight Dynamics team point of view, the hardest challenges consisted in the difficulty of the navigation tasks, the frequent commanding cycles with short turn-around time and the complex setup and trouble-shooting of the TGO GNC autonomous functions, all of which required continuous analysis work and refinement of tools and procedures.

This paper presented some of the most relevant aspects of the aerobraking navigation problem, how they were faced by the ESOC Flight Dynamics team and what results were achieved. The steep learning curve permitted to collect important lessons and identify critical points.

The first aspects regard the spacecraft design: the aerobraking technique requires an agile satellite with a low ballistic coefficient and a high thermal insulation. Agility, intended as fast slew capability, allows to reduce the time spent in attitude manoeuvres and thus to increase the margins around atmospheric passes and/or the duration of Earth tracking periods for navigation and commanding. Both these aspects are critical for the robustness of aerobraking operations, especially for the shortest orbit periods. Low ballistic coefficients and a high thermal insulation permit instead to increase the ΔV per pass and shorten the overall duration of operations, reducing the required man-power and the overall operational risk. Another design aspect,

related to spacecraft autonomy, is the importance of propagating the PTE-based timeline shift to the execution of autonomous manoeuvres: this was not available on TGO and would have drastically reduced the requirements on the on-ground orbit predictions, extending the commanding horizon and alleviating the need of frequent navigation cycles.

Regarding navigation approach and methods, one key aspect can be highlighted: despite the efforts spent in implementing and validating complex atmospheric models, TGO aerobraking confirmed the high unpredictability of the Martian atmosphere. The entire navigation system shall therefore be designed to be robust against sudden density variations, adopting the appropriate margins in the design of the aerobraking corridor. The chosen margin of 150% proved to be adequate, and represented a good compromise between mission safety and duration of operations. Moreover, the adoption of the simple exponential density model resulted in an easier interpretation of the density variations, compared to the more sophisticated MCD and MarsGRAM: in spite of slightly better statistics in the description of the individual passes, none of these models guaranteed a significant improvement in orbit predictability. Orbit predictions during aerobraking thus remain a critical and difficult task, and at the current state no alternative approach has been identified than running navigation cycles with very high frequency.

Other important navigation considerations derive from the analysis of the single triggering of an autonomous protection manoeuvre. This was caused by a combination of a fast walk-in strategy, which did not leave enough time to gather atmospheric density data before designing the next lowering manoeuvre, and an inadequate value of the scale height parameter of the density model. The minor violation of the heat flux limit had no consequence, but it could have been avoided with a closer monitoring of this fundamental quantity and an earlier update of the model, together with a slower walk-in procedure.

The Doppler-based orbit determination system proved very reliable, but required extensive ground station coverage, especially for shorter orbit periods. Relatively simple dynamic models are in general sufficient to reconstruct the orbit and gather information on the atmosphere, but the inclusion of telemetry-based information eased the OD setup and improved the accuracy of the results. In particular the use of accelerometer measurements proved valuable and is highly recommended, either directly as dynamic model or as observations in the OD. It can in fact help the estimation of all components of the drag force and reduce the demand of ground station time, allowing to distinguish the effect of the individual braking passes during tracking data gaps.

As a final comment, the authors would like to stress the very demanding requirements in term of man-power imposed by aerobraking: even in presence of spacecraft autonomy, a large and well trained operations team is necessary, capable of running critical tasks over extended working shifts for a prolonged period of time, without reaching exhaustion. This should be carefully considered in the design and planning phase of any mission employing the aerobraking technique.

Acknowledgments

The authors would like to acknowledge all the colleagues in the ESOC Flight Dynamics team for interplanetary missions, as well as the Flight Control Team and all other ESA support teams, who contributed to the success of TGO's aerobraking campaign. Special thanks to Ana Mestre, Vishnu Janarthanan, David Jesch, Alejandro López, Bernard Godard and Trevor Morley for their substantial contributions to the work presented in this paper. A special mention goes also to the TGO navigation team at JPL, for the solid and patient operational support and the invaluable exchange of expertise.

References

1. Svedhem, H., and O’Flaherty, K., “Exomars 2016 Mission, Brief Description of TGO and Schiaparelli”, ESA public report, January 2016.
2. Yamaguchi, T., and Kahn, M., “Mission Design and Navigation Analysis of the ESA Exomars Program”, *24th International Symposium on Space Flight Dynamics (ISSFD)*, Laurel, Maryland, USA, May 2014.
3. Cassi, C., et al., “The ExoMars 2016 Mission Flight Performances until Achievement of the 1-sol Orbit”, *68th International Astronautical Congress (IAC)*, Adelaide, Australia, September 2017.
4. Denis, M., et al., “Thousand Times through the Atmosphere of Mars: Aerobraking the ExoMars Trace Gas Orbiter”, *SpaceOps 2018*, Marseille, France, May 2018.
5. Castellini, F., Bellei, G. and Godard, G., “Flight Dynamics Operational Experience from ExoMars TGO aerobraking campaign at Mars”, *SpaceOps 2018*, Marseille, France, May 2018.
6. Giorgini, J., et al., “Magellan Aerobrake Navigation”, *British Interplanetary Society, Journal*, vol. 48, no. 3, March 1995.
7. Esposito, P., et al., “Navigating Mars Global Surveyor Through The Martian Atmosphere: Aerobraking 2”, AAS 99-443, *AAS/AIAA Astrodynamics Specialist Conference*, Girdwood, Alaska, August 1999.
8. Johnston, M.D., et al., “The Strategy for the Second Phase of Aerobraking Mars Global Surveyor”, AAS 99-303, *AAS/AIAA Astrodynamics Specialist Conference*, Girdwood, Alaska, August 1999.
9. Smith, J.C., and Bell, J.L., “2001 Mars Odyssey Aerobraking”, *Journal of Spacecraft and Rockets*, Vol. 42, No. 3, May-June 2005.
10. Mase, R.A., Antreasian, P.G., Bell, J.L., Martin-Mur, T.J., and Smith J.C., “The Mars Odyssey Navigation Experience”, AIAA 2002-4530, *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Monterey, California, August 2002.
11. You, T.H., et al., “Navigating Mars Reconnaissance Orbiter: Launch through Primary Science Orbit”, AIAA 2008-6093, *AIAA SPACE 2007 Conference & Exposition*, Long Beach, California, September 2007.
12. Long, S.M., et al., “Mars Reconnaissance Orbiter Aerobraking Navigation Operation”, AIAA 2008-3349, *SpaceOps 2008*, Heidelberg, Germany, May 2008.
13. Damiani, S., Garcia, J.M., Guilanya, R., Munoz, P., and Muller, M., “Flight Dynamics Operations for Venus Express Aerobraking Campaign: a Successful End of Life Experiment”, *25th International Symposium on Space Flight Dynamics ISSFD*, Munich, Germany, October 2015.
14. Shane, R.W., and Tolson, R.H., “Aerothermodynamics of the Mars Global Surveyor Spacecraft”, *NASA technical report NASA/CR-1998-206941*, March 1998.
15. Renault, H., et al., “Exomars 2016, Trace Gas Orbiter GNC Design and In-Flight GNC Behaviour”, *GNC 2017*, Salzburg, Austria, June 2017.
16. Renault, H., et al., “ExoMars 2016, Orbiter Module Bus a GNC Development Update”, *CEAS Space Journal*, Volume 7, Issue 2, pp 105–118, 2015.
17. Guilanya, R., and Companys, V., “Operational Approach for the Exomars Aerobraking,” *23rd International Symposium on Space Flight Dynamics (ISSFD)*, Pasadena, California, USA, November 2012.
18. Tarzi, Z.B., Berry, D.S., and Roncoli, R.B., “An Updated Process for Automated Deepspace Conjunction Assessment”, *25th International Symposium on Space Flight Dynamics ISSFD*, Munich, Germany, October 2015.

19. Long, S., et al. "Mars Reconnaissance Orbiter Aerobraking Daily Operations And Collision Avoidance." *20th International Symposium on Space Flight Dynamics ISSFD*, Annapolis, Maryland USA, September 2007.
20. Konopliv, A. S., Park, R. S., Folkner, W.M., "An Improved JPL Mars Gravity Field and Orientation from Mars Orbiter and Lander Tracking Data", *Icarus*, Volume 274, Pages 253-260, 2016.
21. Bierman, G. J., "Factorization Methods for Discrete Sequential Estimation", *Academic Press*, 1977.
22. Budnik, F. and Mackenzie, R. A. , "An Advanced Modular Facility for Orbit Determination of ESA's Interplanetary Missions", *14th International Symposium on Space Flight Dynamics ISSFD*, Los Angeles, CA, 2001.
23. Budnik, F., Morley, T. A., and Mackenzie, R. A. , "ESOC's System for Interplanetary Orbit Determination", *16th International Symposium on Space Flight Dynamics ISSFD*, Munich, Germany, 2004.
24. Forget, F., Millour, E., and Lewis, S. R., "Mars Climate Database V5.2 Detailed Design Document", CNES, February 2015.
25. Justh, H. L., "Mars Global Reference Atmospheric Model 2010 Version: Users Guide", NASA/TM-2014-217499, Marshall Space Flight Center, Huntsville, Alabama, February 2014.
26. Young, B., "Practical Orbit Determination for Aerobraking with Accumulated Accelerometer Data", *John Junkins Dynamical Systems Symposium*, College Station, Texas, May 2018.
27. Tolson, R. H., et al., "Application of Accelerometer Data to Atmospheric Modeling During Mars Aerobraking Operations", *Journal of Spacecraft and Rockets*, Vol. 44, No. 6, November–December 2007.