Refactoring the AOCS/GNC Sub-System from HERA to COMET-INTERCEPTOR - a Lesson in Minimizing Validation Efforts Using Small-Body Heritage

Francisco Cabral(1), Pablo Colmenarejo(2), José Francisco Briz(2), Pedro Cachim(1), Maciej Quoos(3), Oleksandr Naronin(3), Julia Wajoras(3), Jorge Arnedo(2), João Pedro Rei(1), Simon Chalkley(4), Claire McCrorie(4), Kieran Hayward(4), Carlos Corral Van Damme(5), Manuel Sanchez Gestido(5), Andrew Hyslop(5)

(1) GMV, Alameda dos Oceanos 115, 1990-392 Lisbon, Portugal, +351 21 382 93 66, {francisco.cabral,pedro.cachim,joao.rei}@gmv.com
(2) GMV, Isaac Newton 11, 28760 Tres Cantos, Madrid, Spain, +34 91 807 21 00, {pcolmena,jbriz,jarnedo}@gmv.com
(3) GMV, Ul. Hrubieszowska 2, 01-209 Warsaw, Poland, +48 223 955 165, {mquoos,anaronin,jwajoras}@gmv.com
(4) TAS, Thales Alenia Space, Building 660, Bristol Business Park, Coldharbour Lane, BS16 1EJ, Bristol, United Kingdom, +44 (0) 117 9375228, {simon.chalkley,claire.mccrorie,kieran.hayward}@thalesaleniaspace.com
(5) ESA ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands, +31 (0)71 565 6565, {carlos.corral.van.damme,manuel.sanchez.gestido,andrew.hyslop}@esa.int

ABSTRACT

This paper describes the process of re-using GMV’s past heritage on small-body GNC/AOCS systems, in close cooperation with our team’s Prime Thales Alenia Space UK and our client, ESA, to adapt it to a new mission during phase A/B – Comet Interceptor. The specific needs of this mission were assessed and compared to GMV’s past developments – notably HERA – to assess the possible re-use of the GNC/AOCS modes and functions. Critically, the impact on the simulation and validation chain is also assessed since Comet-Interceptor, as an F-class mission, prioritises solutions with high implementation speed and cost-efficiency for the implementation and validation of its sub-systems. Required new developments are also identified and are discussed in this paper as well. In the end, a subset of the performance tests is presented to showcase the performance of the developed system.

1 Introduction

Space explorations missions, especially smaller class small-body missions, have been recently prioritizing time- and cost-efficient implementations. The easiest approach is to start from an existing sub-system and re-factor it to the new mission.

In the scope of the Comet Interceptor (Comet-I) mission, we discuss here such an approach for the AOGNC sub-system (the junction of the AOCS and GNC systems) with a focus on the GNC component during the closest phases of the comet fly-by – this was identified as the most challenging phase of the mission.

The re-use of the system is most critical on the architecture and mode design – these are the designs that most affect the validation effort and associated required time. Thus, their re-use already guarantees a great leap forward in efficiency. Indeed, these were re-used to a high degree from HERA to Comet-I with only small changes to adapt it to the new mission’s context. The algorithms,

on the other hand, are easier to modify with smaller impact on the validation effort – but these changes shall also be minimized in as much as possible.
This work presents an account of this re-use analysis and identification and design of the required new developments for the context of the Comet-I (when compared to the HERA mission). This focuses around three key functions (image processing, navigation filter, and on-board camera exposure control) and one key validation tool – the comet image generator (CIG).
We start with a review of the Comet-Interceptor (section 2) followed by the design of the GNC strategy for the fly-by phase (section 3). The core of the work is documented in section 4 with the identification of the key design drivers (4.1), followed by a re-use analysis (4.2) and the account of the design of new developments (4.3). The performance test campaign results are provided in section 4.4. The conclusions will close this work in section 5.

2 Comet-Interceptor Mission Overview

2.1 Mission Profile

Comet Interceptor is a fast-track development mission (F-Mission) aiming to perform multi-point observations of a Dynamically New Comet (DNC) or a long-period comet (LPC) from a close flyby with, among other remote-sensing and in-situ instruments, its very narrow FoV scientific camera – CoCa (0.69deg x 0.92deg).
The mission will be launched with an Ariane 62 from Kourou to the Sun-Earth Lagrangian point L2 (SEL2). It will share the launch with the ESA mission ARIEL in a standard dual launch configuration, with Comet Interceptor on top of the dual launch structure, and ARIEL inside it. The planned launch date is by the end of 2029.
Following a waiting phase at the Sun-Earth L2 point of up to 3 years used to select the actual target comet, Comet Interceptor will start its transfer phase lasting up to 4 years. It will then perform a transfer manoeuvre to depart from L2 to intercept the target comet (potentially including a swing-by at the Moon for an extra "push").
Given the target will be selected very late in the development (even after launch), the mission and system design has to be performed to cope with a wide range of parameters, notably with fly-by velocities between 10 and 70 km/s, solar aspect angles between 45 and 135 degrees and distances to the Sun between 0.9 and 1.2 AU. This is a significant challenge for the design of the AOGNC sub-system.
After the cruise phase and a few weeks before the Encounter phase the SC will detect the target comet and perform the required guidance manoeuvres. On the last 3 days before the closest approach point, the SC will enter its critical operations phase with ground entering its most intensive phase to command the SC to perform the required manoeuvres to guide the spacecraft (SC) to deliver two small probes approximately 24h and 48h before the closest approach (C/A), probes B1 and B2, respectively. After their release, the two probes will perform autonomous operations, relaying the scientific data back to the main spacecraft.
Following the release of the probes, the SC will perform a manoeuvre to guide the SC to its final trajectory with a C/A point 1000km from the comet. After this manoeuvre the control of the SC is handed over to the autonomous GNC that will point the scientific instruments to the target's body while protecting the S/C from the incoming coma dust with a shield installed in the leading face of the S/C.
Following the encounter, the main spacecraft will downlink to earth the data acquired during the fly-by.

2.2 Comet-I and HERA Comparison

Comet-Interceptor is still a small-body mission but one that has key differences to the HERA

mission. Thus, under the approach to re-use the HERA GNC S/S, it is important to highlight the key differences and similarities between both missions – this will feed into the re-use analysis in 4.2 for the AOGNC S/S.
At the levels of the key similarities, we identified the needs for an autonomous attitude guidance functionality (based on an on-board optical navigation filter), for a similar AOGNC HW sensor suite, and having similar operational, safety and robustness requirements. These, together with the need for a cost-efficient and fast development and validation campaign for this sub-system has driven the decision to re-factor the sub-system from the HERA mission [9].
At the same time, it is critical to identify the most important differences to ensure a proper re-factoring of the subsystem:

- The target will only be selected shortly before launch or in flight thus, the AOGNC must be designed for a range of possible parameters. The range of parameters is as wide as:
  - Velocities between 10 and 70 km/s
  - Solar aspect angles between 45 and 135 degrees
  - Closest approach distances between 550 and 1450km (3σ limits, considering the navigation errors around the nominal closest approach distance of 1000 km)
  - Heliocentric distances between 0.9 and 1.2AU.
  - Unknown target size, shape, physical and optical properties.

  This presents specific challenges to the GNC in order to ensure a) correct exposure of the target on the navigation images and b) estimate the location of the centre of the target.
- The comet will have a coma which is expected to be significant enough to create challenges in the image processing and star tracker performances (with expected outage periods). It also will be composed of several microscopic dust particles that, at very high-speeds, will create torques on the spacecraft and can impact the surface of the optical instruments lenses affecting their performance. A shield must be mounted on the SC to protect against these dust impacts that creates significant attitude constraints that the AOGNC must follow.
- The target must be captured in a very narrow FoV payload (0.92x0.69deg) resulting in tight APE (absolute performance error) requirements during the Encounter phase (APE smaller than 190mdeg in the flight direction). The payload will be mounted in a rotating mirror (discussed in 3) and, consequently, the AOGNC must be able to effectively communicate with the rotating mechanism and provide the required operating inputs.

This sets up the context for the AOGNC strategy and subsystem designs.

3 The GNC Fly-by Strategy

After the last divert manoeuvre, 20h before the closest approach (C/A), the SC is left pointing to a 1000km closest-approach passage (with a 150km 1σ dispersion). This marks the last ground update and the start of the encounter phase. After the last ground contact, the spacecraft's state must be computed by the on-board GNC to correctly command the SC's attitude. The payload command inputs are also computed to correctly point the narrow FoV of the CoCa instrument (0.69deg x 0.92deg) by providing the required angular profiles to the rotating mechanism in which this instrument is mounted on.
One of the main challenges of the encounter phase is the coma cruising. During the time the SC is inside this coma, it will suffer the impact of very small dust particles that, due to the high fly-by velocities can damage the platform and/or introduce significant torques that the AOCS must compensate or recover from. The number and magnitude of these impacts increases as the SC gets closer to the comet's nucleus. In order to protect the spacecraft, a partly-fixed attitude during the coma cruising is maintained and, this way, only the platform's leading face (or ram face) is shielded

(Figure 1); the SC only maintains the freedom to rotate about its velocity direction.

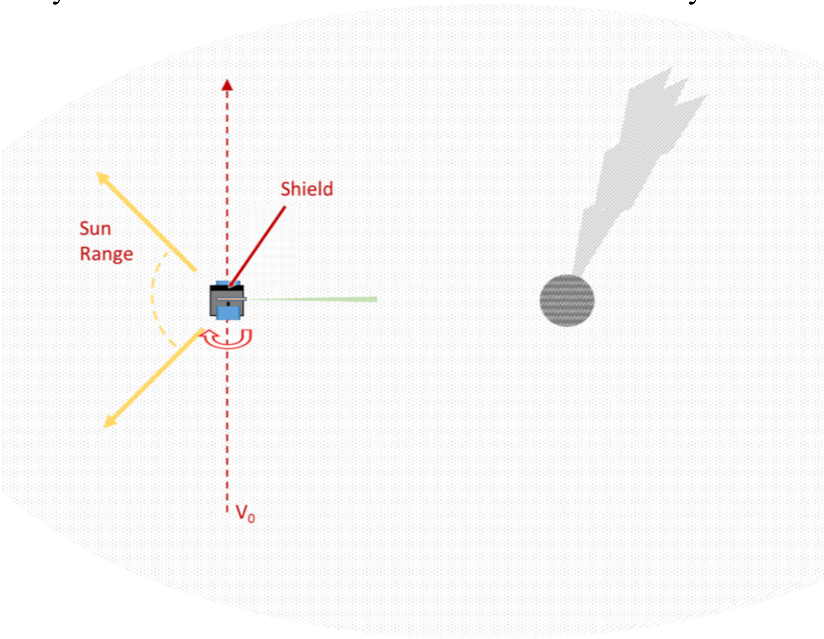

Figure 1: Encounter Strategy: the target maintains a fixed attitude to have only its ram face shielded – the SC can only rotate around its ram direction.

Under this strategy, in order for the CoCa to point at the target, a rotating mirror assembly (RMA, Figure 2) is mounted in the SC to allow the instrument rotation about one axis. This axis is maintained (by the SC commanded attitude) to be normal to the SC's flight plane. This way, the RMA can be commanded to point the CoCa towards the target during its flight path.

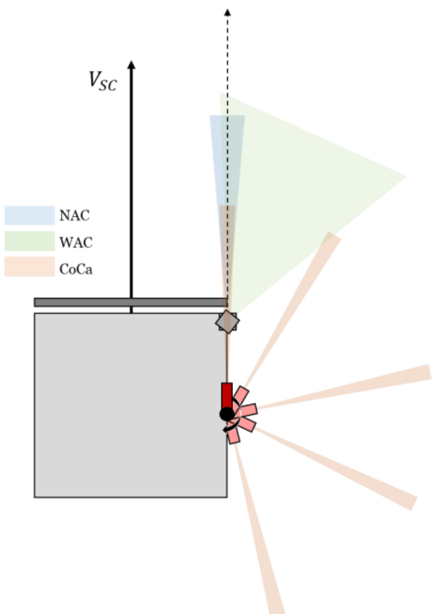

Figure 2: Rotating Mirror Assembly (RMA) allowing different orientations of the CoCa.

The proposed SC design during the competitive phase A/B would have two nominal WACs on-board, each with a FoV of 68x54deg, to observe the target, perform on-board navigation and provide the required information to command the SC's attitude and the CoCa and MIRMIS payload on-board the SC. The configuration of the WACs was the subject to a trade-off and the selected approach was to have the first WAC mounted in such a way that it has the ram direction in its FoV, just close to the edge (Figure 3). A second WAC expands this FoV to span a total > 90deg to allow observation of the target from the beginning of the phase all the way to the C/A point (remember that the SC cannot rotate in that direction to maintain a shield-first attitude).

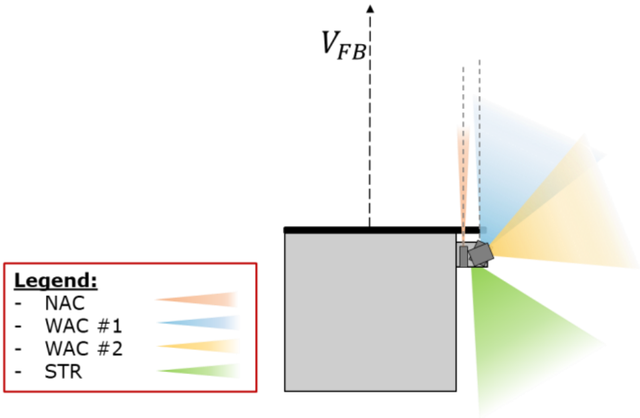

Figure 3: Proposed WAC Configuration (Phase A/B).

Other configurations were tested – to not include the ram direction (not observing the target at the beginning) or having a gap between the WACs – but these resulted in configurations with issues at the robustness level on the image processing, filter, and autonomous exposure control performances. Indeed, having constant observation results in a more robust approach allowing ground to have a verification of the GNC performance and for the GNC to constantly track the target's brightness levels ensuring a proper exposure during the critical phase close to the C/A. A continuous observation also allows more robust reactions to any failures.

It is important to mention that even though there will be a NAC on-board, this is not used by the GNC due to cost constraints. The purpose of the NAC is to provide the required accuracy for ground to achieve, during the approach phase, the required payload delivery performance by spacecrafts A, B1 and B2.

At some point, the target will eventually be lost from the WAC's FoV (Figure 6) and the GNC will then perform pure propagation to obtain the SC's position solution to command the SC's attitude and payload command inputs.

In summary, the GNC has three key responsibilities during the Encounter phase: They are to provide command inputs for:

- The SC attitude about its dust shield normal;
- The CoCa to point to the target on its motion about 1 axis;
- The MIRMIS (another payload with a wider FoV) to point to the target on its motion about 1 axis.

To respect the system constraints, the SC attitude must be commanded such that:

- The dust shield normal vector is aligned with the velocity vector;
- The RMA rotation axis is perpendicular to the spacecraft A's flight plane (relative to the target); this gives two solutions.

The ambiguity is broken by the target on the side covered by the RMA.

3.1 The Challenge of Imaging a Comet

For a mission based on on-board optical navigation in the proximity of the target, it is critical to understand the appearance of the target. Thus, we analyse the expected target observation conditions. These type of comets – LPCs or DNCs – are fairly unknown bodies so all expected behaviour should be taken as, at best, an educated guess. We highlight [8] for a specialised survey on the expected appearance of comets.

One would expect the target to be, with a high probability but not certainly, a high activity target (as it must be detected several years in advance from Earth while Comet-I is stationed at L2) and to

have one or several jets with between mild to strong activity. Most jets should originate on the illuminated side of the target with higher probability and intensity on the terminator and local noon. The jets direction is roughly normal to the surface and thus, not necessarily on the Sun's direction.
In fact, there are situations where the direction of the Sun and that of the jet are opposite on the camera sensor (see Figure 5). Note that this is not an unlikely occurrence, and it should occur at some point for the fly-by's approaching with a solar aspect angle below 90 degrees. Thus, we reach the important conclusion that the jet direction cannot be inferred from the sun direction in the camera detector plane – this is an important point for the design of the image processing function.
During the 20h where the SC will observe the target, the comet's nucleus will be smaller than a pixel (Figure 4) but its coma will be bright and large enough that the camera can easily detect it. At around 2.5-20 minutes before the C/A the target will become resolved (depends mostly on the fly-by velocity and target radius) and the jets will become distinguishable. Even closer to the C/A, the SC's position will rotate fast around the comet's nucleus. This mean that the observed side of the target will also change fast with potentially new jets coming into the camera's view and past ones disappearing. This will be another challenge to the image processing as the location of the brightest pixel or region might suffer large jumps due to this rotation of jets coming/disappearing from the camera's field of view.

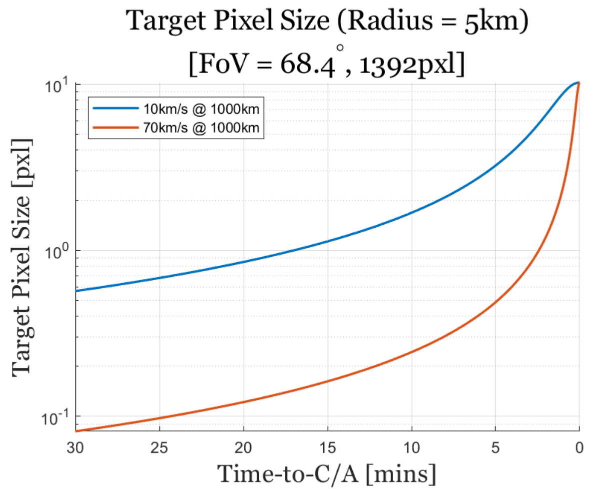

Figure 4: Comet's body size in pixels with respect to time to go for the fly-by limit velocities.

There is yet another possible factor for a fast change of the brightest region in the field-of-view of the camera. If the SC crosses the jet's direction (see Figure 6), which is a likely occurrence for any fly-by with SAA < 90 deg, so about half of the cases, the brightest region will suffer a fast change from one side of the target to the other. This can happen also if the jet anti-direction is crossed (SAA close to 135deg)
Indeed, it has been identified in the simulation results that this can be the driving source of mis-performances in some simulation shots. Thus, the fast change of the brightest region in the target in the camera detector frame is the most challenging observation characteristic that must be accounted for in the design of the image processing function.
To deal with the potential fast changes of the brightest spot GMV has implemented a bias estimation by a Gauss-Markov process as suggested on NASA's best practices for on-board navigation [5]. This strategy is discussed in more detail in section 4.3.4.
GMV has also made a design surrounding more conservative assumption such as, possibility of brightest spot away from the limb, high jet-to-coma ratios, and outburst occurrences. This, however, are more unlikely scenarios but make part of a conservative and robust design.
It is also important to note that, although the high activity of the comet is an expected characteristic it is not a certain one. Thus, the GNC image processing function must be designed to deliver high performances for both high and low-activity targets.

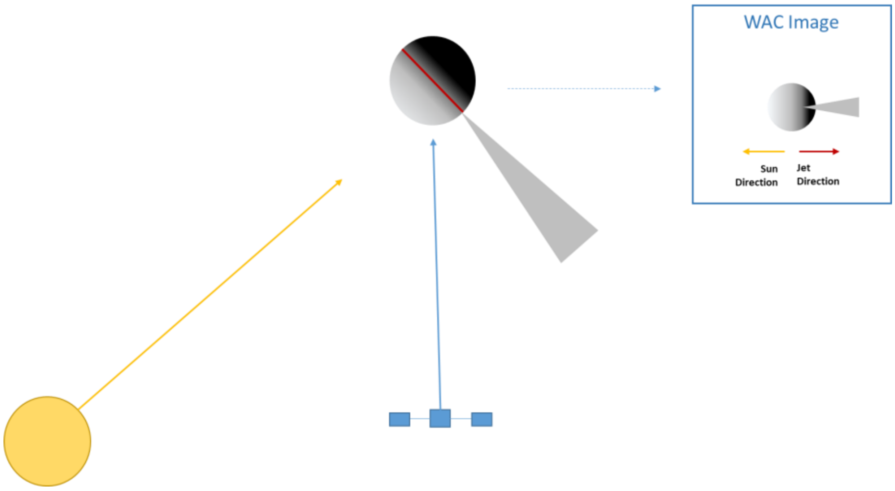

Figure 5: Jet direction and Sun direction, when decomposed into the camera detector plane, do not have to match.

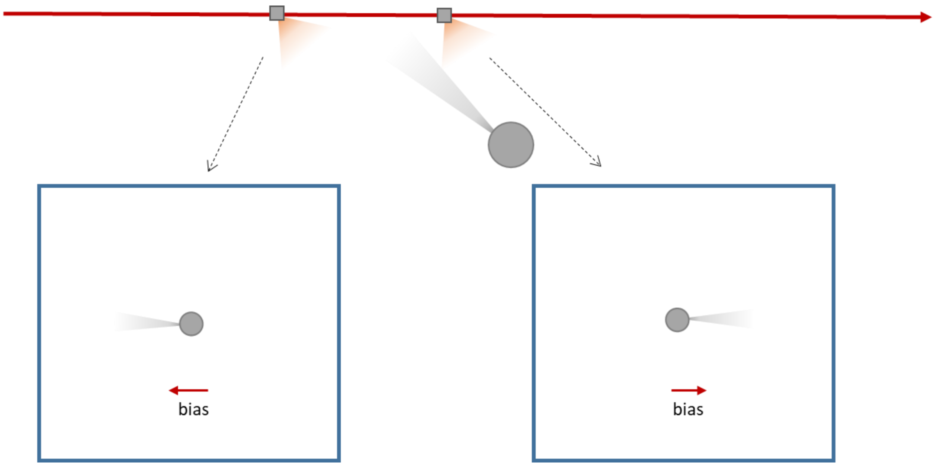

Figure 6: Expected sudden change on the jet direction on the detector as a consequence of the geometry of the fly-by. This is an expected occurrence for SAA < 90deg.

4 Re-Factoring Approach for the Encounter Approach Mode

Once the general GNC strategy is defined the identification of the key design drivers can be defined. This is followed by a re-use analysis of the HERA architecture and functions and the design of the required new developments. This account of our approach in the Comet-I is finalised with a small subset of the results obtained in the test campaign.

4.1 Identification of the Key Design Drivers

The key design drivers guided the decision on which functions are re-used from HERA and which are developed specifically for the Comet-Interceptor GNC. There were three key identified drivers:

- The GNC must be robust to different target and coma sizes, shapes, and optical properties.
- "Hit-or-Miss" scenario: during the Encounter, the GNC must possesses extra robustness in its nominal performances and optimize performance under outages of some sensors such as the star trackers or potential temporary unavailability of navigation images. During this phase safe mode transitions are blocked to avoid going into safe mode, since this is not acceptable in this phase)
- Excellent navigation performance under challenging conditions to ensure correct pointing of the CoCa payload.

4.2 Re-use Analysis

The next step in the design of the COMET-I GNC function is the identification of the HERA's GNC architectures and functionalities that can be re-used for Comet-I and which require new developments – this is convened, appropriately, the re-use analysis:

- **GNC Mode Architecture:** The mode architecture is almost all maintained from HERA; a few modes that are not needed in the scope of Comet-I are discarded. There are, however, changes at the sub-mode level to reflect the different algorithmic needs of Comet-I as well as the different conditions within the Encounter phase.
- **GNC Functional Architecture:** this is kept from HERA which was already heritage from previous developments. GMV has a very flexible and adaptable background functional architecture that is used within most its operational developments, and this was re-used for Comet-I.
- **Image Processing:** the very different nature of a comet and its coma and the wide range of target properties that the IP function must be able to cope with required a new development for this function.
- **Translational Navigation:** HERA's translational navigation algorithm would be suitable for Comet-Interceptor but it was decided, due to the nature of the observability conditions, that there were some available techniques that could enhance the navigation performance for Comet-I. Thus, some formulation within the algorithm were modified to try to meet the mission's ambitious performance requirements. However, the navigations function architecture and interfaces were re-used from HERA.
- **Attitude Guidance:** HERA's attitude guidance architecture and functionality was re-used but a final step was introduced to separate the required attitude into a "shield first" attitude for the SC with a rotation around the ram direction and an angular profile for the rotating mechanism of the main scientific payload – CoCa. This was a simple modification. An attitude smoother at the guidance output was also implemented to ensure that the inputs sent to the rotating mirror did not violate actuation constraints (otherwise the commands would require the rotating mirror to do small jumps between solutions).
- **GNC-FDI:** the FDI functionalities from HERA's GNC are re-used with one significant difference – in a "hit-or-miss" scenario, sending the spacecraft into any non-nominal modes could dictate failure in achieving the mission objectives linked to the comet observation so alternative recovery actions were designed.

In the next section we address the key new developments with emphasis on the image processing, and comet image generation capabilities for the GNC functional engineering simulator (FES).

4.3 New Developments

In this section we focus on the three new key developments that were performed specifically for COMET-I – the comet image generator (CIG), the image processing, the autonomous camera exposure control, and the on-board navigation filter.

4.3.1 Comet Image Generator (CIG)

At the start of the study of the GNC S/S for Comet-I, there was a study performed on how the target for this mission could look like. The closest target in terms of characteristics was the Halley comet, observed during the fly-by from the long past Giotto Mission in 1986 (see Figure 7). As one can imagine, engineers did not have, at the time, the tools for documenting and archiving data that we do have today and the amount of the available data in this mission was limited. But even Halley was a short period comet while Comet-I is targeting a long-period comet. Other comet missions had targets that were significantly less active than Halley so extrapolating from their behaviour should be performed with extreme care.


Figure 7: Giotto approaching Comet Halley [3].

After this first review of what the target could be expected to look like, research followed to assess both the available tools to simulate navigation images during the Encounter and, in an alternative approach, the available platforms to develop a new comet image generator tool. Available tools like PANGU, ASTOS camera simulator, and IGOR (an in-house development in openGL) were assessed first but they were found to not have the embedded capabilities to properly simulate a comet. The development of a new tool was then the selected approach.
For the platform to perform this new development, several renderers were studied: Blender, 3DS Max, Unreal Engine, and more. To select the tool best fit for our purposes a trade-off was performed with different criteria – ease of use, quality, runtime performance, possibility of interfacing to the GNC FES in Simulink, license cost – which culminated in the selection of Blender to develop this key functionality for the FES.
A full account of this development is an extensive topic that merits a stand-alone account. Nevertheless, we summarise here the main features and capabilities implemented in the tool developed based on Blender that was named Comet Image Generator – CIG:

- **Nucleus** Modelling
  - The nucleus size is defined as an input that is used to set the overall shape mean radius.
  - The nucleus shape irregularity was modelled irregularly by adding Perlin noises at different frequencies in a stochastic manner.
- **Coma** Modelling
  - The coma was modelled as a cloud of dust around the nucleus.
  - The coma shape was slightly elliptic, elongated towards the Sun's direction.
- **Jets** Modelling
  - Comet jets were modelled as originating on the illuminated face of the nucleus.
  - The number and base strength of the jets can be defined, each with different apertures resulting in different intensities and maximum (identifiable) size simulated in a stochastic way.
  - The probability of the origin of the jet is higher on the terminators and especially close to the local noon of the nucleus. Their origin is then defined stochastic with a defined

probability distribution.
    - The modelling of the jets is performed by volumetrics as opposed to the coma to ensure a different ratio of strength between the jets and coma depending on the distance of the observer – at further distances the coma dominates the brightness of the scene opposed to closer distances where this brightness is dominated by the jets. Furthermore, the jets strength falls off as an almost inverse square of distance from the comet as opposed to the coma that falls in an inverse linearly relation with the distance. Table 1 provides a matrix of images generated with different jets/comas strength ratios before calibration and background.
    - The jets are not uniform and have a Musgrave noise texture implemented in a stochastic way to emulate local variability.
- **Outburst** Modelling
    - Outburst are modelled as a strong narrow jet that can randomly appear on nucleus surface and last for seconds to minutes.
    - Outburst also increases the total coma brightness.
    - Their modelling is similar to that of jets (but different parameters such as much larger strength).
- **Sun Illumination** Modelling
    - Provided directly by Blender
    - Important to note that the selected method to model each comet's feature was selected as to make them look realistic under different illumination conditions.
    - It assumes that the Sun provides parallel rays, i.e., that it is in the infinite.
- **Stars** Modelling
    - Stars are generated randomly.
    - They could be generated from a catalogue but this capability was not included at this stage as it affected runtime and the stars are not being used in the image for navigation but we still wish to see how they affect the processing of the comet.
- **Scene Brightness**
    - The overall scene brightness changes depending on an input for the exposure time of the image. This can result in overexposed or underexposed images.
    - This is modelled with parameters that define the camera sensor gains and comet parameters. This is critical to assess the autonomous camera exposure control function.
- **Camera** Modelling
    - FoV and number of pixels are modelled in the image along with camera effects such as lens distortion and sensor noise.

We note that there was an intensive calibration of the jets/models parameters to ensure that the output resulted in realistic images at different distances to the comet when the different parameters were dispersed. Without this calibration, dispersing the strength of one feature without modifying the parameters of the other could result in images with unrealistic behaviour (e.g.: jets dominating at very far distances and coma at close distances or fall-off behaviours between coma and jets that were unrealistic).

Table 1: Different CIG Images with different coma and jets strengths before calibration and without stars or camera errors to showcase the variability possibility within the tool.

	Coma density multiplier 0.3	**Coma density multiplier 0.8**
Jets density multiplier 0.01	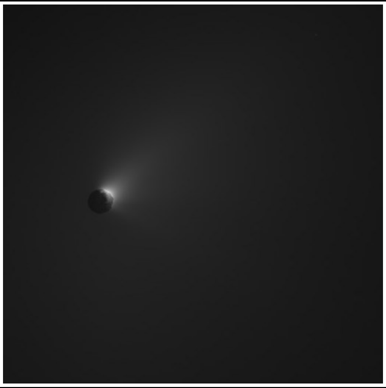	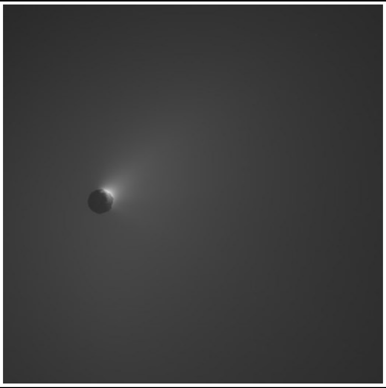
Jets density multiplier 0.05	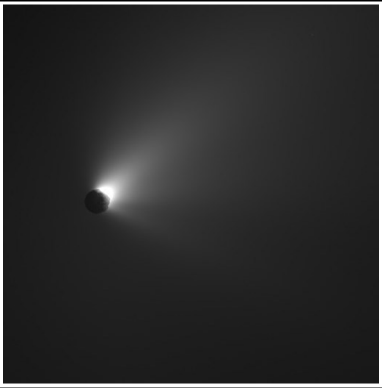	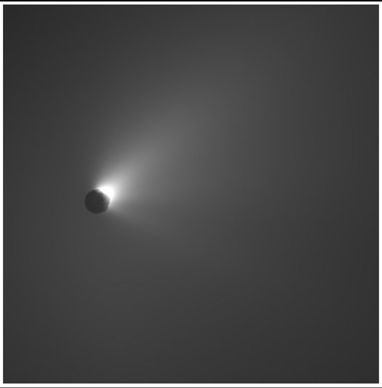

4.3.2 Image Processing

The image processing performance was found to be the number 1 performance driver in the GNC Encounter phase during the sensitivity analysis performed early in the development phase. Thus, it soon become an area of development focus: Three different algorithms were tested:

- **HotPix** Algorithm: takes the *n* brightest pixels (*n* is usually taken < 10pixels) and performs centroiding on them. This is then taken as the nucleus centre. Due to its simplicity this algorithm is very robust.
- **Jet-Compensation** Algorithm: this algorithm singles out the jet from the rest of the image (blob identification) and determines its direction by a geometrical fit. This direction is then used to identify the jets' origin that is used to identify the nucleus centre. This algorithm presents better performance than the HotPix but lesser robustness. In this algorithm, the HotPix was also run in the background so that if their offset distance was larger than twice the radius of the target (in pixels) the measurement was discarded in favour of the HotPix. Figure 8 shows this algorithm in action.
- **Hybrid** Algorithm: Another approach is to mix the above two approaches. It was found that when the jet compensation algorithm fitted the single out jet to a shape close to a circle, then either the coma dominated the scene brightness or the jet had its direction in line with the LoS (line of sight) of the SC – in both situations the HotPix presented better performance. The ratio between the axes of the fitted shape to the jet was then used as the variable in a switching function where a ratio of 1 led to exclusive use of the HotPix and a ratio of 0.5 led to the exclusive use of the Jet Compensation algorithm.

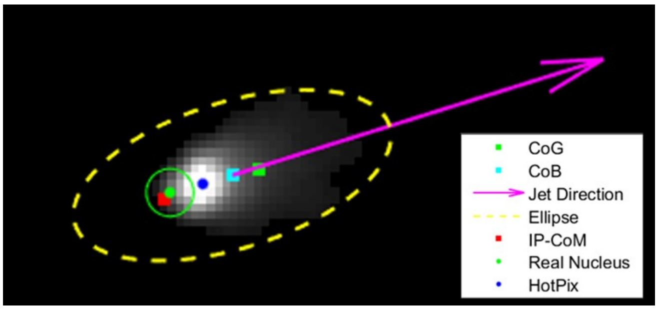

Figure 8: Jet Compensation Algorithm. In yellow the shape used to determine the jet's direction (in magenta) which is then used to determine the origin of the jet.

In any trade-off of IP techniques, one has to be mindful of the dangers of over-fitting, i.e., designing an algorithm that performs very well under a set of parameters or with a specific image generator tool but not in real life when the images are slightly different than expected. In order to avoid overfitting an extensive sensitivity analysis and stress tests were performed which included:

- Different jets/coma strength ratios (on different extremes: jets dominating, coma dominating, both dominating at different distances);
- Brightest pixel in the jet is away from the nucleus;
- Jet not directed towards the Sun (this led to the decision of not using any information on the Sun's direction in the IP algorithms);
- Dominating jet directed towards the SC;
- Dominating jet directed away from the SC;
- Outbursts;
- Underexposed Images;
- Overexposed Images.

Following these tests, the Hybrid algorithm was selected as it presented the best compromise between performance and robustness. However, we note that the simulation campaign documented in 4.4 was still run with the HotPix as a conservative approach to estimate the expected GNC performances.
There was one stress case that severely impacted the IP performance – overexposed images. This led to the decision to develop a function for the autonomous exposure control. The IP provides outputs that feeds directly into that function – the pixel DN (data number) histogram and number of saturated pixels. Note that the IP pre-processing already performed a first clean on the image of any hot pixels or artifacts.

4.3.3 Autonomous Camera Exposure Control

The performance of the IP, critical to the overall GNC performance during the encounter, is dependent on receiving well exposed images. In the context of this mission this is no easy task – the target may present a high brightness variability, especially at closer distances where phenomena like jets and outbursts gain a higher predominance/probability. Thus, a functionality to autonomously control the camera exposure time was introduced.
This functionality will receive inputs from the IP and, keeping a history of the latest image values, it will control the exposure over the next cycles. The used controller is a scheduled PID controller whose scheduling parameter is the distance between subsequent shots. The I and D gains are proportional to the P gain to reduce the degrees of freedom of the problem
The controller is defined in the logarithmic space since the exposure affects the DNs through multiplication. The logarithmic of the exposure is taken and the PID output is added to that. The exponent of the resulting value is the new exposure that is fed to the camera.
The target for the controller is the average of the n (set by parameter) brightest pixels inside the

blob identified by the image processing. It is also possible to discard the *m* brightest pixel in this consideration. This strategy is used to avoid being impacted by bright stars or artifacts that might not have been discarded by the IP. The precise value of *m* and *n* is defined by an absolute value and a relative value that varied with the size of the target's nucleus in pixels (and thus varies with the proximity to the target).
An initial guess functionality was also implemented that guaranteed fast convergence at the beginning of the encounter phase or if, for example, a temporary error led to the loss of navigations images for a few minutes during which the scene brightness would have changed significantly.
The authors note that this functionality was yet not run in the performance test reported in section 4.4 but it was validated in an open loop validation test.

4.3.4 On-Board Navigation Filter

The architecture of the on-board navigation filter functionality is the same as in HERA. However, the Comet-I Encounter phase presented some key challenges that led us to a re-evaluation of the used algorithm:

- Observability depends on the distance and angle to the target and this changes very rapidly close to the target;
- In certain situations, the IP estimate errors are highly non-linear and can have "jumps" that can lead to an overestimation of the SC's velocity (e.g.: SC crossing the jet direction).
- "Loss-of-target" could result in a later recovery with an a priori solution that is much worse than the nominal performance. This can lead to a known error in sequential filters where the large improvement of the filter update can lead to linearization errors creeping in the gain function when considering the "high" value of the residuals.

HERA first used an Unscented Kalman Filter (UKF) – Extended Kalman Filter (EKF) hybrid (the unscented transformation as only applied to the observables step and not to the propagation step for the a priori solution). Later, this was converted into a full EKF (the unscented transformation did not provide significant performance in the HERA scenario). This simple yet robust filter provided all the required performance in HERA.
For Comet-I, we decided to use the sequential UDU Information filter [2] with a Gauss-Markov process to model the observation noise and an underweighting factor [5][6]. This uses four formulations that are often seen in navigation implementations and are of great advantage for our scenario:

- The **UD decomposition** increases the filter robustness by avoiding negative eigenvalues in the covariance matrices due to numerical errors when the process noise is relatively small and the eigenvalues are close to zero. This formulation accomplishes this by avoiding the square root operations performed in the Kalman filter. This is particularly relevant when relative large updates (high gains) can present from the rapid change in observability close to the C/A in the fly-by.
- The **information formulation** allows to set up no a priori information (information matrix equal to zero; the covariance matrix cannot be set to infinite in the traditional formulation) and allows to avoid the inversion of the covariance matrix, an operation that is expensive and introduces numerical errors – this is another advantage of the information filter that, similarly to the UD formulation, reduces the chances of filter divergence due to new high precision observations.
- When a bias is present in the measurements or dynamics, it can be included in the state vector so it is estimated and improves the general performance of the other affected state parameters. This is the case in Comet-I where there will be an error between the IP measurement and the actual nucleus centre that is correlated on time. However, biases without dynamical evolutions are usually locked in very fast by the filter and, because biases

are seldomly pure biases, the filter can easily lock in to a solution that is either off or will be off in the near future. This can lead to filter underperformance or, in the worst cases, filter divergence. The solution is to estimate biases by a **Gauss-Markov process** [5]. This still presents an improvement in the overall filter behaviour by estimation of the bias but, with a dynamical behaviour, it avoids locking into a solution and might present updates if the dynamics and/or observables behaviour demand it. This technique increases significantly performance and robustness of the filter. In Comet-I this is implement only for the observation function. No significant bias on the dynamical behaviour with real impact on the filter performance was identified.

- In non-nominal situations, when the target is lost from the FoV of the camera momentarily (either due to an anomalous unexpected situation, sensor error, large dust particle impact, or other system error), the filter might have to perform a very fast improvement of the solution once measurements are recovered – this is problematic for the filter. The filter will improve its covariance very fast (because the new measurements are much better) but the a priori solution is off from the real solution and thus significant non-linearity errors creep in the filter gain. This can lead to filter divergence as the filter locks in a wrong solution too fast. A solution to this problem is to implement an **underweighting factor** [5] – this prevents the filter to correct too much in one step and, thus, non-linearities are corrected little by little as the new measurements' improvement is also introduced progressively. There are different ways to implement this factor (see [6] for a thorough discussion on this) but, for Comet-I, the inverse of the trace of R (measurement covariance) has found to be the best approach. Using a dynamic underweighting factor based on the relationship between the dynamics second-order effects and the measurement expected errors (HPH^T) has proven to have a slightly better performance but not enough to justify the added computational complexity.

It is important to note that the implementation of the Gauss-Markov process requires tuning of the parameters defining the process (time correlation factor and spectral density) which is non-trivial.

4.4 Performance Validation Campaign

A MonteCarlo (MC) performance test campaign was performed to validate the performance of the overall GNC Encounter performance. The number of MC shots was selected to achieve a 90% confidence for the intended requirement: a minimum of a 95% success rate according to the ECSS standards [7] for the requested APE performance – 190mdeg in the camera frame in the flight direction and 290 mdeg in the perpendicular direction over 99.7% of the time within a +/-2hr window.

4.4.1 Model-In-The-Loop (MIL) Set-Up

The nominal scenario only assumed 1 STR and presented STR blinding when the coma become too bright around the C/A. Dust particle impacts with impact on the APE error were also part of the nominal scenario.
The number of dispersed parameters in the MC campaign was quite extensive:

- Initial SC and target heliocentric state (knowledge and dispersion error)
- Encounter fly-by velocity, solar aspect angle, C/A point
- Target radius, activity, optical parameters (entering CIG) and mass
- SC dynamics errors (SRP factor, non-gravitational acceleration)
- Target dynamics errors (non-gravitational acceleration)
- SC platform parameters (MCI parameters, CoM offset)
- Sensor errors (mounting errors and instrument errors)
- Actuators errors (mounting errors and actuator errors). This includes rotating mirror actuation errors that have a direct impact on the performance that GNC cannot influence.

A high number of sensitivity tests was also run to a) fully understand the performance drivers and b) help construct the pointing budget. These are not reported here but included: no dust particle impacts, IP performance models (white noise only, bias only and, critically, a single step from 3σ to 3σ margins; the latter provided to be, by far, the worst case), no STR blinding, no mounting misalignment errors, WAC #1 failure, WAC #2 failure, and a RW failure scenario.

4.4.2 Hardware-In-The-Loop (HIL) Set-Up

The HIL-OPT introduces an actual HW camera to produce the images that the IP will process, and the CIG generated images are displayed on a LED monitor. Mechanical components are used to ensure the proper relative mounting of camera and display. The HIL elements are contained in GMV's optical laboratory (***opt-lab***) to prevent light disturbances to affect images.

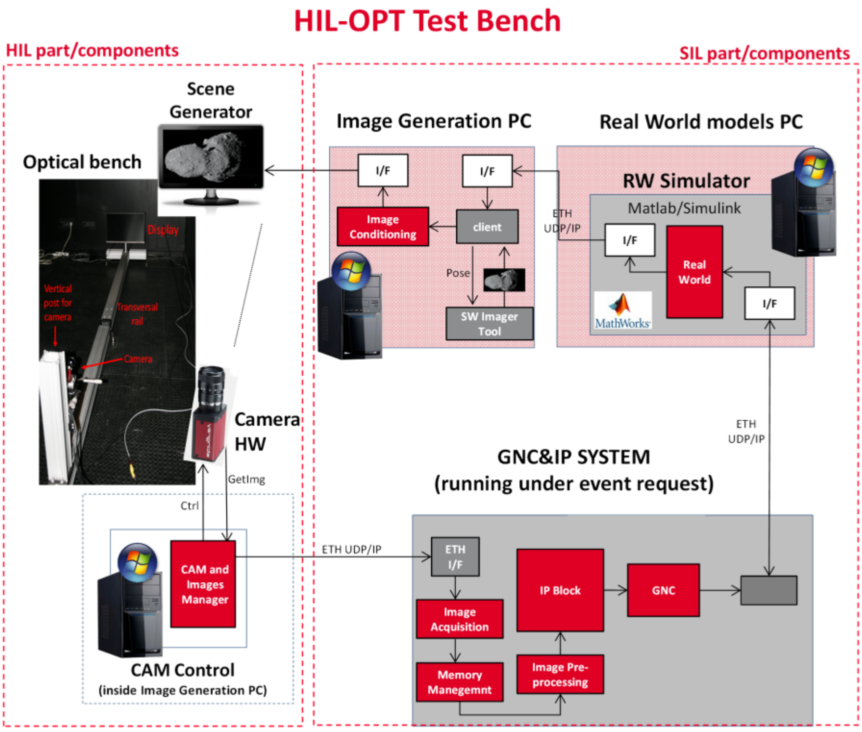

Figure 9: HIL-OPT Architecture.

- The SIL components are those in the right part of Figure 9:
  - The Real-World (RW) Simulator containing the FES Real world models.
  - The Image Generation PC in charge of generating representative images generated with the CIG and including image conditioning and interfaces both with the RW simulator PC and with the scene generator screen.
  - The tested GNC SW resides in the GNC&IP System, hosted on the same PC running the RW and Image Generation.
- The HIL elements are those in the left part of the figure.

  1. **Scene Display**:

  - Electronic display (screen), where the generated images are displayed. From a preliminary selection the monitor used is a LED-PLS screen of 24'' diagonal and 3840X2160-pixel resolution.
  - The SW function running in the Image Generation PC is in charge of (1) conditioning (as needed) the image created by the CIG, (2) feeding the stimulation screen with the image, and (3) controlling the display window to stimulate correctly the working area of the camera

detector.
- The HW interface in the Image Generation PC in charge of linking with the electronic display at physical level.
- The entire scene generator is mounted within a “dark room” - the GMV’s Optical Navigation Laboratory - to prevent the degradation of captured images by unfavourable ambient illumination.

2. **Camera HW**:

- The camera body is the AVT Manta G-419-B. This commercial camera is preliminary selected as (1) its technology (CMOS) and resolution (2048×2048 pixels) are representative of typical space navigation cameras, (2) its data interface (Giga Ethernet) permits a high frame rate (28.6 fps at full resolution), (3) its mechanical interface with the lens (C-mount) offers the widest availability of commercial objectives (focal length, aperture, quality), (4) it is SW controllable in exposure and gain, with the subsequent flexibility of use, and (5) the detector temperature is part of the camera telemetry, thus allowing the image conditioning SW to accurately remove image bias.
- Camera lens: the choice of the camera lens takes into consideration (1) the camera (detector size), (2) the scenario (field of view), and (3) desired quality factors of the lens (resolution). Lens selection shall be performed (as customization) for the different scenario phases.

3. **CAM Control PC**: this component is implemented inside the same PC of the Image Generation PC. It has the following tasks:

- Camera control, monitoring and images downloading. These functionalities are implemented via the camera manufacturer SDK libraries.
- Image calibration function in charge of applying the required calibration and formatting to the raw image retrieved from the camera.
- Communication bus with the GNC/IP Computer.

4. **Optical Bench:** the mechanical train that ensures that camera and display are mounted and positioned in a steady manner and can be accurately adjusted to the relative pose that every test requires.

Figure 10 presents an image with the MANTA camera and screen mounted in the optical bench during the calibration process. It is important to note the criticality of this calibration process as it has significant impact on the overall test performance. This calibration is performed at 4 critical levels: alignment (centre, roll, and perpendicular); edge cropping, photometric calibration and, finally, geometric calibration to correct any lens and projective distortions.

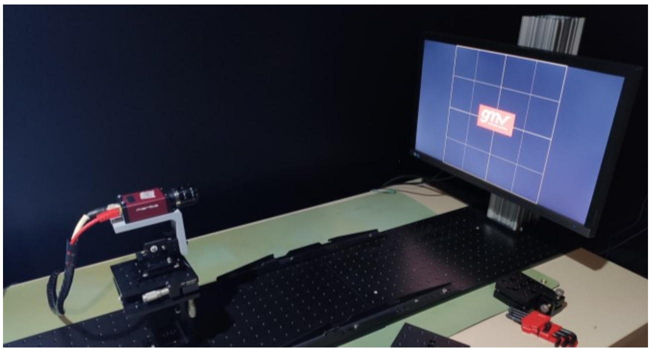

Figure 10: Optical Bench with the mounted camera and screen.

4.4.3 Results

The results of the MIL and HIL campaign are presented in Table 2 and Figure 11. It is important to note that the requirement allows a total of 0.3% of the time outside the performance threshold (~40s). This time is used almost entirely recovering from dust particle impacts that are much more frequent at closer distances.

Another important note is the fact that the autonomous exposure control function was still not in the loop in this set of simulations. Its integration within the GNC loop was foreseen for the next iteration loop as the final calibration of the CIG tool was only finalized in the end of the last iteration loop (already included in these tests).

Table 2: Performance Results of MIL and HIL campaigns

	Cases Fully Compliant with the CoCa APE Requirement
MIL	99%
HIL	99%

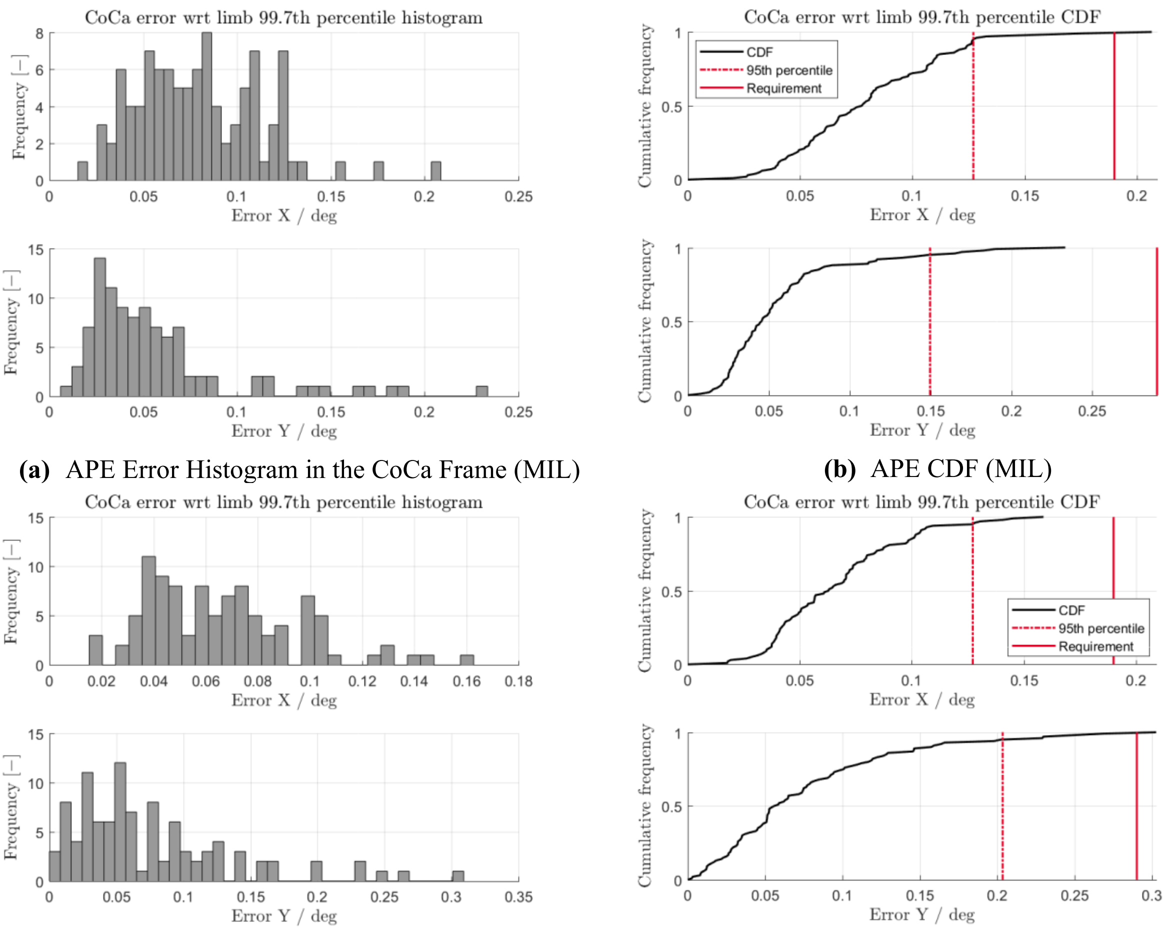

(a) APE Error Histogram in the CoCa Frame (MIL) **(b)** APE CDF (MIL)

(c) APE Error Histogram in the CoCa Frame (HIL) **(d)** APE CDF (HIL)

Figure 11: APE error in the CoCa frame and percentage of successful case over each step. A successful case is one that is within the threshold 99.7% of the time within a +/-2h window.

The MIL and HIL simulation campaign were run using the HotPix IP algorithm as a conservative approach. The implementation of the hybrid IP algorithm was foreseen for a later stage which are expected to bring performance improvements.

5 Conclusions

We have provided an account of the TAS-GMV consortium work on the development of the AOGNC S/S focused around the GNC modes for the Encounter phase. This was the most challenging phase of the mission.

This was a fast and cost-efficient development that challenged the limits of the available technologies. The development of the AOGNC S/S was performed in a fast 18 months that span phases A to B2 and achieved TRL 5-6 with a very limited budget/effort.

Even if the TAS-GMV consortium was not selected for the implementation phase, the performed work and efforts are presented in this paper as a feasibility proof of efficient AOGNC re-use/update from one deep-space mission to another and for a possible future re-use. Given the flexibility of re-factoring the system from HERA to Comet-I, and the important new developments performed here that are very relevant to any comet or fly-by mission, this is a product that can be easily adapted to future missions.

The work was not fully finalised and a small number of items remained to be fully closed (that were expected to phase C). In specific: the repeat of the of the HIL simulations with the selected camera (whose selections was only closed by the end of B2), the implementation of the hybrid IP algorithm into the test campaign to collect their associated performance gains, and the inclusion of the autonomous exposure control in the performance test campaign.

6 Acknowledgements

The authors wish to acknowledge all ESA and TAS personnel that also provided reviews, comments and iterations in the AOGNC sub-system design. This work is also thanks to them.

7 References

[1] ESA, *Comet Interceptor: Consolidated Report on Mission Analysis*, version 1.1, December 2020

[2] Zanetti, Renato & D'Souza, Christopher, *Information Formulation of the UDU Kalman Filter*. IEEE Transactions on Aerospace and Electronic Systems. 55. 493–498, 2019, 10.1109/TAES.2018.2850379.

[3] ESA website, *ESA - Giotto approaching Comet Halley*, https://www.esa.int/ESA_Multimedia/Images/2016/03/Giotto_approaching_Comet_Halley

[4] Tapley, B., Schutz, B., Born, G., *Statistical Orbit Determination*, Elsevier Academic Press, 2004

[5] Carpenter, J.R., D'Souza, C. N., *NASA Best Practices*, First Edition, April 2018

[6] Zanetti, R., DeMars, K. J., and Bishop, R. H. *Underweighting nonlinear measurements*. Journal of Guidance, Control, and Dynamics 33, 5, 1670–1675, 2010

[7] ECSS standards, *Space Engineering: Control Performance Guidelines*, Handbook, ECSS-E-HB-60-10A, Dec 2010.

[8] Vincent, JB., Farnham, T., Kührt, E. et al. *Local Manifestations of Cometary Activity*. Space Sci Rev 215, 30, 2019.

[9] Palomino, A., et. al., *Design, Development, Validation and Verification of the HERA GNC subsystem*, ESA GNC Conference 2023, June 12th-16th, Sopot, Poland.