



HSC

Doc. no.: HERSCHEL-HSC-TN-2069
Issue: 2.1
Date: 20 October 2016
Page: 1

Guide to the Herschel Gyro-based Attitude Reconstruction Software

Craig Stephenson

20 October 2016



Contents

1	Introduction	3
2	Method	5
2.1	Improvement of star tracker attitude measurements	5
2.1.1	On-board construction of measured star vectors	6
2.1.2	On-ground processing	8
2.2	Gyro-based attitude reconstruction	17
2.2.1	Derivation of the model	17
2.2.2	Estimation of the model parameters	20
2.2.3	Calculation of the goodness-of-fit	23
2.2.4	Computation of the new attitude estimate and its uncertainty	24
3	How to use the software	26
4	Summary of known issues	47
A	Inversion of distortion correction equations	49
A.1	Using a modified Newton method	49
A.2	Using power series	50
B	Ignoring the inter-axis correlation of the measurement errors	52
C	Normality of the measurement errors	57



1 Introduction

Despite the success which was achieved during the mission in improving the accuracy of the star tracker attitude measurements, first on-board [7, 23, 24] and later on-ground [8], it has been shown that the on-board attitude filter (upon which the ground-based pointing reconstruction was also based) is rather poor at following the high-frequency changes in the spacecraft attitude (i.e. the spacecraft jitter) [2, 3].¹ A new method, the ‘gyro-based attitude reconstruction’ method, was therefore proposed for combining the star tracker attitude measurements with the output from the Inertial Reference Unit (GYR) [9]. As the name implies, this method places much greater weight on the measurements made by the gyros, using attitude measurements constructed from star tracker data to provide an absolute reference and to account for gyro drifts.

Having demonstrated itself to be capable of successfully reducing the high-frequency components of the Absolute Measurement Error (AME), the gyro-based method has become the default attitude reconstruction method in the pipeline software—its predecessor (the application of a simple correction to the attitude estimated on-board) being now only used under exceptional circumstances. Nevertheless, the method is not perfect and has its own limitations, which might have been avoided by the implementation of a more conventional attitude estimator.² However, there is no evidence that these limitations are of any real importance or have compromised the resulting attitude estimates.³

The gyro-based attitude reconstruction method was first prototyped in September 2012 by Helmut Feuchtgruber [9] and subsequently implemented

¹The consequence of improved star tracker attitude measurements is a reduction in the bias and low-frequency components of the Absolute Measurement Error (AME). When the AME is reduced on-board it has the additional effect of simultaneously reducing the Absolute Pointing Error (APE).

²In August 2013, a more general-purpose estimator (in which low-pass filtered attitude measurements from the star tracker are combined with high-pass filtered gyro measurements) was prototyped and shown to produce equally-accurate results for the case of staring observations [24]. Mark Tuttlebee further showed that, by simply retuning the gains, the on-board filter may be used to produce similar results.

³One limitation is that, when ‘calibrating’ the gyro measurements, only those attitude measurements close to the reference attitude may be used. Another is that a deterministic model of the variation of the gyro drift rates must be assumed (the gyro drift rates are assumed to remain constant over intervals of a fixed length).



in Jython, as part of the pointing toolbox, by Bart Vandenbussche (March 2013). My own involvement began in April 2013 during the preparation of the software for its first release in HIPE 12.0 (June 2014) and I have since continued to make further refinements to both the method and its implementation.⁴ These refinements—many of which were quite substantial—were incorporated into HIPE 13.0 (March 2015) and HIPE 14.0 (December 2015). Towards the end of 2014, Jonathan Cook rewrote the software in Java so that it could be incorporated into the HCSS auxiliary processors. All changes which were since made to the Jython version of the software—apart from those related to the generation of diagnostic output—were also implemented in the Java version. As of HIPE 15.0, the pointing toolbox invokes the same software used within the auxiliary processors, the original Jython code being now no longer used.⁵

This document serves two main purposes: to provide a detailed description of the gyro-based method (Section 2); and to describe how to use the standalone version of the software (Section 3). A list of the outstanding issues known to be affecting the latest build of the software is also included (Section 4).

⁴Shortly prior to the release of HIPE 12.0, a version of the software which I had until then been using for testing became the official version. Since that date any errors introduced into the software can be blamed on the author of this document!

⁵A consequence of this is that the generation of diagnostic output (and the selection of the redundant gyro) is no longer available in HIPE 15.



2 Method

2.1 Improvement of star tracker attitude measurements

Although not truly part of the gyro-based attitude reconstruction method, the first function of the software is to use data contained in the STR-specific diagnostics telemetry dataset (`AcmsDtmStr`), whenever these data are present, to create a set of improved (often referred to as ‘corrected’) star tracker attitude measurements.⁶ For each timestep (typically every 1 s, but every 0.25 s is also possible) and for each of the stars tracked by the star tracker, the following information is used:

- y_b - and z_b -coordinates, `strnnPosY` and `strnnPosZ`, of the aberration-corrected, measured star vector in the Boresight Reference Frame (BRF). In what follows we will refer to these coordinates as u'_{y_b} and u'_{z_b} .
- Catalogue ID, `strnncatId`, of the matched star;
- right ascension and declination, `strnnRa` and `strnnDec`, of the matched star (corrected for proper motion);
- the parameter α_c , `alpha_c_nn`. This is used to correct the star tracker focal length according to the colour of the star; see (4).

Typically, the above information is present for up to 9 stars, i.e. $nn \in \{01, 02, \dots, 09\}$, but for observations for which interlacing mode was active there may be as many as 18 star vector measurements, i.e. $nn \in \{01, 02, \dots, 18\}$.⁷ In the latter case, any measurements for which $nn \in \{10, 11, \dots, 18\}$ came from the previous ACMS cycle and need to be handled slightly differently (see Section 2.1.2).

Using these data and the residual distortion maps [see 8], the software constructs improved measurements of the star directions and these in turn

⁶In this document the term ‘star tracker attitude measurement’ refers to an estimate of the spacecraft attitude, i.e. of the Attitude Control Axes (ACA) reference frame, that is made using data downlinked from the star tracker. It does not refer to a measurement of the star tracker attitude, i.e. of the Boresight Reference Frame (BRF).

⁷Unfortunately, due to an error in the definition of version MOC-12 of the Mission Information Base (MIB), although star tracker interlacing mode was used on-board from very early in the mission (from at least OD 200), the data associated with this mode did not become available in telemetry until OD 866 (see HCSS-20596).



are used to create improved attitude measurements (replacing those that were originally produced on-board by the star tracker).

Before explaining in detail how the improved attitude measurements are created, it is helpful to review the procedure employed on-board the spacecraft to convert the measured CCD coordinates of each star into an aberration-corrected, measured star vector, \mathbf{u}' .

2.1.1 On-board construction of measured star vectors

The process which was used on-board by the star tracker to convert the measured CCD coordinates of each star into a geometric (i.e. free from the effects of stellar aberration), measured star vector can be broken down into three steps:

1. *Correction for star tracker CCD distortion*

Let the position of the star on the CCD, as determined by the centroiding algorithm, be (y, z) .⁸ The distortion-corrected coordinates, (y', z') , were calculated on-board according to:⁹

$$\begin{aligned} y' &= F(y, z; k_0, \dots, k_7), \\ z' &= F(z, y; h_0, \dots, h_7), \end{aligned} \tag{1}$$

where

$$\begin{aligned} F(u, v; a_0, \dots, a_7) &= a_0 + a_1u + a_2v + a_3u(u^2 + v^2) + a_4u(u^2 + v^2)^2 \\ &\quad + a_5u^2 + a_6uv + a_7v^2, \end{aligned} \tag{2}$$

and the a_i ($i = 0, \dots, 7$) are real constants. The coefficients k_0, k_1, \dots, k_7 and h_0, h_1, \dots, h_7 are referred to as the ‘distortion correction coefficients’ for the y - and the z -axis respectively.¹⁰

⁸The variables y and z correspond respectively to the variables y_{fcd} and z_{fcd} in [4, p. 83].

⁹Eqs. (1) and (2) correspond precisely to eqs. (9.2-3) in [4, p. 83], with y' and z' replacing y and z , and y and z replacing y_{fcd} and z_{fcd} .

¹⁰Three sets of values for the parameters k_i and h_i were used on-board during the mission, depending on the Operational Day (OD): (i) $OD \in \{1, 2, \dots, 865\} \setminus \{858\}$: pre-launch values; (ii) $OD \in \{858\} \cup \{866, 867, \dots, 1010\} \setminus \{1005\}$: modified values of k_1 and h_1 ; (iii) $OD \in \{1005\} \cup \{1011, 1012, \dots, 1446\}$: all values modified.



2. Calculation of the focal length and apparent star direction

The apparent direction of the star with respect to the BRF is given by the unit vector:¹¹

$$\mathbf{u} = \begin{pmatrix} f \\ -y' \\ -z' \end{pmatrix} (f^2 + y'^2 + z'^2)^{-\frac{1}{2}}, \quad (3)$$

where the focal length, f , of the star tracker is determined from:¹²

$$f \equiv f(f_0, \alpha_c) = f_0[1 + \alpha_T(T - T_0) + \alpha_c]. \quad (4)$$

f_0 represents the nominal focal length of the star tracker at the reference temperature, $T_0 = 22$ C, and the parameter α_c (as obtained from the entry of the matched star in the on-board star catalogue) is used to correct the nominal focal length according to the colour of the star.¹³ For the entire duration of the mission, the coefficient of thermal expansion of the objective, α_T , was set equal to 5.91×10^{-6} C⁻¹ and the temperature, T , of the star tracker optics (which may found in the housekeeping telemetry) was found to remain constant at 13.73 C.¹⁴

3. Removal of the stellar aberration

Stellar aberration is removed using the classical model [6]. That is, the (unit) vector in the true (geometric) direction of the star is given by:

$$\mathbf{u}' = \mathbf{u} + \mathbf{u} \times (\mathbf{u} \times \mathbf{v}_{s/c})/c, \quad (5)$$

where $\mathbf{v}_{s/c}$ is the spacecraft velocity vector and c is the speed of light.¹⁵

Whenever interlacing mode is active there will be star vector measurements associated with both the current ACMS cycle and the previous ACMS cycle. Before using the measurements to determine the attitude of the star tracker, the estimated spacecraft rates are used to propagate (i.e. rotate) those measurements which were made at the previous cycle to the time associated with the current frame [see 1, pp. 190–191].

¹¹Eq. (3) corresponds precisely to eqs. (9.2-5) in [4, p. 83], with y' and z' replacing y and z .

¹²See eq. (9.2-1) of [4, p. 82].

¹³Two values for the nominal focal length, f_0 , were used on-board during the mission: $f_0 = 29.96330$ mm (OD < 762), $f_0 = 29.97178$ mm (otherwise).

¹⁴See [8, p. 6].

¹⁵Since \mathbf{u}' is required with respect to the BRF, the spacecraft velocity vector, $\mathbf{v}_{s/c}$, is transformed from inertial coordinates to BRF before using eq. (5).



2.1.2 On-ground processing

The process by which the attitude reconstruction software constructs, at each timestep (i.e. every 1 s), an improved attitude measurement is the following:

Removal of the on-board distortion correction

The first part of the process is to take each of the (aberration-free) measured star vectors

$$\mathbf{u}' = \begin{pmatrix} \sqrt{1 - (u'_{y_b})^2 - (u'_{z_b})^2} \\ u'_{y_b} \\ u'_{z_b} \end{pmatrix}, \quad (6)$$

(expressed here in the BRF) and to remove the distortion correction that was applied on-board, thereby recovering the measured position of the star on the CCD. The detailed steps required are:

1. If interlacing mode is active and the measurement belongs to the set associated with the previous ACMS cycle, then rotate \mathbf{u}' back to its (BRF) direction one ACMS cycle earlier.¹⁶ That is, if \mathbf{u}'_t represents the vector at time t , then:

$$\mathbf{u}'_{t-\Delta t} = \mathbf{q}_r(\Delta t) \mathbf{u}'_t \mathbf{q}_r^*(\Delta t), \quad (7)$$

where

$$\mathbf{q}_r(\Delta t) \approx \mathbf{q}_{\text{align}} \left[\frac{\hat{\omega}_{\text{aca}} \Delta t}{2} \right] \mathbf{q}_{\text{align}}^*, \quad (8)$$

$\mathbf{q}_{\text{align}}$ is the star tracker alignment quaternion,¹⁷ $\hat{\omega}_{\text{aca}}$ is the estimated

¹⁶Interlacing mode is considered to be active whenever the flag `strmInterlacStatus` in the main 4Hz ACMS dataset `AcmsScmTM` is non-zero. It is by default assumed that it is those star measurements in positions 10–18 which are associated with the previous ACMS cycle:

In interlaced mode, the star information . . . of stars measured in the last cycle are *presented before* the star information measured in the previous cycle and updated to the current one. [4, p. 38] and [5, p. 36]

¹⁷For most of the mission—other than immediately following a star tracker reset—the operational star tracker was STR-A. The alignment quaternion for this star tracker is: $\mathbf{q}_{\text{align}} = [0.1953842 \times 10^{-3} \quad -0.2993422 \times 10^{-1} \quad -0.9995519 \quad 0.2061553 \times 10^{-4}]^T$, which amounts to a rotation of approximately 180° through an axis close to $z_{\text{str-a}}$ [1,



spacecraft angular velocity vector in the ACA-frame and $\Delta t = 0.25$ s.¹⁸ (In (7) and (8) all vectors are expressed, in the usual way, as quaternions.)

2. Include the stellar aberration to obtain the apparent star direction:

$$\mathbf{u} = \mathbf{u}' - \mathbf{u}' \times (\mathbf{u}' \times \mathbf{v}_{s/c})/c. \quad (9)$$

3. Calculate the distortion-corrected CCD coordinates of the star:¹⁹

$$\begin{aligned} y'_r &= f \frac{u_y}{u_x}, \\ z'_r &= f \frac{u_z}{u_x}, \end{aligned} \quad (10)$$

where u_x , u_y and u_z are the three components of \mathbf{u} expressed in the BRF, and the focal length, f , is calculated from (4) using the value of f_0 that was on-board during the OD in question (see footnote 13) and the value of α_c contained in `AcmsDtmStr`.²⁰

4. Solve

$$\begin{aligned} y'_r &= F_1(y_r, z_r; k_0, \dots, k_7), \\ z'_r &= F_1(z_r, y_r; h_0, \dots, h_7), \end{aligned} \quad (11)$$

where

$$\begin{aligned} F_1(u, v; a_0, \dots, a_7) &= -F(-u, -v; a_0, \dots, a_7) \\ &= -a_0 + a_1u + a_2v + a_3u(u^2 + v^2) + a_4u(u^2 + v^2)^2 \\ &\quad - a_5u^2 - a_6uv - a_7v^2, \end{aligned} \quad (12)$$

and the coefficients k_i and h_i are those that were on-board during the OD in question (see footnote 10), to obtain the measured position (y_r, z_r) of the star on the CCD.

p. 77 and p. 602]. Throughout this document, the quaternion $q_1\mathbf{i} + q_2\mathbf{j} + q_3\mathbf{k} + q_4$ will be represented by the 4-vector $[q_1 \ q_2 \ q_3 \ q_4]^T$.

¹⁸The estimated ACA-frame body rates are obtained from `estAngVelX`, `estAngVelY` and `estAngVelZ` of the `AcmsScmTM` dataset.

¹⁹The subscript 'r' has been used to distinguish the variables used in the attitude reconstruction software from those used on-board.

²⁰Note that eqs. (9.3-2) of [4, p. 84] are incorrect; in each a factor $-f$ is required for consistency with eqs. (9.2-5).



The software computes an approximate solution of (11) by means of

$$\begin{aligned} y_r &= F_1(y'_r, z'_r; k'_0, \dots, k'_7), \\ z_r &= F_1(z'_r, y'_r; h'_0, \dots, h'_7), \end{aligned} \tag{13}$$

where the coefficients k'_i and h'_i of these ‘inverse polynomials’ have been computed so as to remove the distortion corrections applied by (11) [see 8, p. 7 and p. 25]. Despite the seeming lack of theoretical justification for using polynomials of this form, it has been demonstrated (and independently verified) that, within the region of the CCD used to track guide stars (a disc of radius 4.07 mm), the errors introduced by this method are all less than 0.05" [8, p. 7]. An alternative approach is described in Appendix A.

Setting $y'_r = -y'$ and $z'_r = -z'$ in (10) we obtain equations which are consistent with (3) and then setting $y_r = -y$, $z_r = -z$ in (11) we recover (1), showing that the CCD coordinates used by the attitude improvement software are of opposite sign to those used in eq. (9.2-3) of [4, p. 83].²¹ Although it is unclear precisely why it was decided to use coordinates of opposite sign,²² this is of no consequence, provided that consistent transformations are used when re-applying the distortion correction (14) and converting back to a unit vector (15).

Application of the new distortion correction

Having removed the on-board distortion correction and recovered the measured position of the star on the CCD, the new distortion correction may be applied. The new correction consists of replacing (11) by:

$$\begin{aligned} y'_r &= F_1(y_r - \widetilde{\Delta y}_r, z_r - \widetilde{\Delta z}_r; k_0, \dots, k_7), \\ z'_r &= F_1(z_r - \widetilde{\Delta z}_r, y_r - \widetilde{\Delta y}_r; h_0, \dots, h_7), \end{aligned} \tag{14}$$

²¹Alternatively, from (3) and (10) we obtain $y'_r = -y'$, $z'_r = -z'$ and then using (11) and (12):

$$\begin{aligned} y' &= F(-y_r, -z_r; k_0, \dots, k_7), \\ z' &= F(-z_r, -y_r; h_0, \dots, h_7), \end{aligned}$$

from which, upon comparison with (1), gives $y_r = -y$, $z_r = -z$.

²²Perhaps some confusion arose from a failure to notice the sign error in (9.3-2) of [4, p. 84] (see footnote 20), followed by an attempt to compensate for this error by changing the sign of all the terms of even degree in the expression for F , i.e. in (2).



where the values of the distortion correction coefficients, $k_0, \dots, k_7, h_0, \dots, h_7$ are those that were on-board between ODs 320–762 (a period which is referred to as the ‘reference period’ [8, p. 7]).²³

The ‘residual distortions’, Δy_r and Δz_r , are obtained from a pair of maps (one for each axis).²⁴ Each map consists of a 5120×5120 array, where each element of the array is identified with a sub-pixel of the CCD (each pixel having been divided uniformly into $10 \times 10 = 100$ sub-pixels).

The residual distortions have been shown to comprise both local (due to the location of the sub-pixel within the pixel) and global phenomena. To reduce the errors present in the local corrections, without unduly affecting the global corrections, the following averaging is employed. Suppose the star is found to lie in a particular sub-pixel of pixel (m, n) , where $0 \leq m, n \leq 511$. Then $\widetilde{\Delta y}$ and $\widetilde{\Delta z}$ are taken to be the median values of the residual distortions at the corresponding sub-pixel of the 121 pixels centred on pixel (m, n) , i.e. from pixels $\{(i, j) : m - 5 \leq i \leq m + 5, n - 5 \leq j \leq n + 5\}$.²⁵

Having applied the new distortion correction, it remains simply to construct the star (unit) vector,

$$\mathbf{u} = \begin{pmatrix} f \\ y'_r \\ z'_r \end{pmatrix} (f^2 + y'^2_r + z'^2_r)^{-\frac{1}{2}}, \tag{15}$$

where the focal length f is calculated using (4), making sure to use the value of f_0 that was on-board during the reference period,²⁶ and remove the stellar aberration using (5).

If the resulting measurement \mathbf{u}' belongs to the set associated with the previous ACMS cycle, then before it can be used in the attitude determination it needs to be rotated to its (BRF) direction one ACMS cycle later.

²³Although a justification for the decision to use this period as the reference period is given in [8, p. 7], the choice was, in fact, quite arbitrary. The only point of importance is that values of the distortion correction coefficients and of f_0 should be the same as those that were used when producing the residual distortion maps.

²⁴There are two such pairs of maps: one for the period prior to OD 320 (when the set-point for the temperature of the star tracker CCD was reduced to -10 C), and one for the remaining ODs. It is assumed that during each of these periods the distortion of the CCD remained constant.

²⁵Note that this is not the method described in [8, p. 12]. Note also that the software includes checks to ensure that the centroid of the star and all the pixels used for this averaging lie within the operational portion (disc) of the CCD plane.

²⁶The decision to use this value was once again quite arbitrary.



That is, using the notation introduced earlier,

$$\mathbf{u}'_t = \mathbf{q}_r^*(\Delta t) \mathbf{u}'_{t-\Delta t} \mathbf{q}_r(\Delta t). \quad (16)$$

Estimation of the star tracker attitude

From the previous two steps we have obtained a set of up to 18 measured star vectors $\{\mathbf{u}'_1, \mathbf{u}'_2, \dots, \mathbf{u}'_n\}$, $n \leq 18$, each in the BRF and each with the new distortion correction applied.²⁷ (In fact, only stars whose centroid lies inside the operational region of the field-of-view are used for the attitude determination. This is a disc, centred on the image of the boresight, of radius 4.07 mm. That is, all star vectors must be within 7.7° of the boresight direction.) The corresponding (expected) inertial directions $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ are then easily derived from information which may be extracted from the star catalogue using the IDs of the matched stars.

The improved estimate of the spacecraft attitude is taken to be the rotation matrix $A \in \text{SO}(3)$, i.e. the three-dimensional orthogonal matrix with determinant 1, which minimizes the loss (cost) function

$$L(A) \equiv \frac{1}{2} \sum_{i=1}^n a_i |\mathbf{u}'_i - A \mathbf{v}_i|^2, \quad n \geq 2, \quad (17)$$

where the a_i are a set of non-negative weights.²⁸ If one adopts the QUEST measurement model [21], that is, if one assumes that each measured star vector, \mathbf{u}'_i , is related to its true direction, $\mathbf{u}'_{i,\text{true}}$, by

$$\mathbf{u}'_i = \mathbf{u}'_{i,\text{true}} + \Delta \mathbf{u}'_i,$$

where the measurement errors, $\Delta \mathbf{u}'_i$, are mutually independent and normally distributed random variables

$$\Delta \mathbf{u}'_i \sim N(\mathbf{0}, R_i),$$

with covariance matrices

$$R_i = \sigma_i^2 (I_{3 \times 3} - \mathbf{u}'_{i,\text{true}} \mathbf{u}'_{i,\text{true}}^T), \quad (18)$$

²⁷It has been found that, for those periods when interlacing mode was active, the rows of the STR-specific diagnostics telemetry dataset `AcmsDtmStr` often contain repeated measurements of the same stars. When this is the case, only the first occurrence (i.e. that corresponding to the current ACMS cycle) is used.

²⁸This problem was first posed, in an unweighted form, by Wahba [25].



and if, furthermore, one chooses the weights such that

$$a_i = \frac{c}{\sigma_i^2}, \tag{19}$$

where c is some arbitrary constant, then it has been shown [see 18] that the attitude matrix, A^* , which minimizes the loss function $L(A)$ is also the maximum likelihood attitude estimate. Equation (18) is equivalent to assuming that the error $\Delta \mathbf{u}'_i$ is perpendicular to $\mathbf{u}'_{i,\text{true}}$ and that it has an axially-symmetric distribution with variance σ_i^2 .

To minimize the loss function, Davenport's q -method is employed. First, $L(A)$ is rewritten [e.g. 14, p. 360] in terms of the gain function g as:

$$L(A) = \lambda_0 - g(A), \tag{20}$$

where $\lambda_0 = \sum_{i=1}^n a_i$,

$$g(A) = \text{tr}(AB^T),$$

and

$$B = \sum_{i=1}^n a_i \mathbf{u}'_i \mathbf{v}'_i{}^T.$$

So, to minimize $L(A)$ we must maximize $g(A)$. In Davenport's method, the loss function is re-written in terms of the attitude quaternion \mathbf{q} , i.e.

$$\tilde{L}(\mathbf{q}) = \lambda_0 - \tilde{g}(\mathbf{q}), \tag{21}$$

where $\tilde{L}(\mathbf{q}) = L(A(\mathbf{q}))$ and $\tilde{g}(\mathbf{q}) = g(A(\mathbf{q}))$, and it is shown that

$$\tilde{g}(\mathbf{q}) = \mathbf{q}^T K \mathbf{q}, \tag{22}$$

where

$$K = \begin{bmatrix} S - I \text{tr}(B) & \mathbf{z} \\ \mathbf{z}^T & \text{tr}(B) \end{bmatrix},$$

$$S = B + B^T \text{ and } \mathbf{z} = \begin{bmatrix} B_{23} - B_{32} \\ B_{31} - B_{13} \\ B_{12} - B_{21} \end{bmatrix}.$$

Applying Lagrange's "method of the undetermined multiplier" to the problem of maximizing (22) subject to the constraint $\mathbf{q}^T \mathbf{q} = 1$, it is easy to see that our optimal attitude estimate, $\hat{\mathbf{q}}$, is an eigenvector of K and from (22) it is also clear that it corresponds to the maximum eigenvalue, λ_{max} .



Having solved the eigenvalue problem and obtained $\hat{\mathbf{q}}$, it remains simply to multiply $\hat{\mathbf{q}}$ by the appropriate star tracker alignment quaternion, $\mathbf{q}_{\text{align}}$, in order to obtain an estimate, $\hat{\mathbf{q}}_{\text{str}}$, of the ACA-frame attitude based on star tracker data. That is,

$$\hat{\mathbf{q}}_{\text{str}} = \hat{\mathbf{q}} \mathbf{q}_{\text{align}}. \quad (23)$$

Calculation of goodness-of-fit and removal of ‘bad measurements’

Now the value of c in (19) is arbitrary and in the pointing reconstruction software we have set $c = (\sum_{i=1}^n 1/\sigma_i^2)^{-1}$ so as to make $\lambda_0 = 1$. Furthermore, we have chosen to weight all the measurements equally, i.e. $\sigma_i^2 = \sigma^2$, so that $a_i = 1/n$ ($i = 1, \dots, n$).²⁹

In [20] Shuster defines the variable TASTE (for the case where $c = 1$, $a_i = 1/\sigma_i^2$) and shows it to be distributed randomly according to a χ^2 -distribution with $2n - 3$ degrees of freedom.³⁰ In [19] he provides a more general definition:

$$\text{TASTE} = \frac{2L(A^*)}{\lambda_0 \sigma_{\text{tot}}^2} = \frac{2(\lambda_0 - \lambda_{\text{max}})}{\lambda_0 \sigma_{\text{tot}}^2}, \quad (24)$$

where $\sigma_{\text{tot}}^2 = (\sum_{i=1}^n 1/\sigma_i^2)^{-1}$, which makes the result independent of the choice for c (or equivalently λ_0).

For the value of c chosen above, we find that

$$\text{TASTE} = \frac{2n \tilde{L}(\hat{\mathbf{q}})}{\sigma^2} = \frac{2n(1 - \lambda_{\text{max}})}{\sigma^2}. \quad (25)$$

The variable TASTE provides an indication of the goodness-of-fit: the larger the value of TASTE, the worse the fit. Assuming our measurement model is realistic, we expect:

$$p_{\text{taste}} \equiv \Pr\{\text{TASTE} > \text{val}\} = 1 - P\left(\frac{2n - 3}{2}, \frac{\text{val}}{2}\right), \quad (26)$$

where P is the lower (regularized) incomplete gamma function. If the value of p_{taste} is less than the value of `prob_thresh` (see p. 29, Table 2), then an

²⁹The estimation of σ is described in the next section.

³⁰This result, which seems very familiar, is normally only true for least-squares problems in which the model is linear with respect to the fitted parameters [16, p. 654]. Wahba’s problem also differs from a standard least-squares problem in that the parameters are constrained by the attitude belonging to $\text{SO}(3)$.



attempt is made to improve the fit by discarding ‘bad’ star measurements. The process consists of removing one bad measurement (the worst) at a time until either p_{taste} exceeds `prob_thresh` or a total of five measurements have been excluded. A measurement is considered bad if (and only if) excluding it from the fit increases the value of p_{taste} by a factor greater than `prob_frac` (see p. 29, Table 2 and footnote 47).

Estimation of the measurement uncertainty

Although when estimating the attitude using a given set of star measurements at time t_i we assume that all measurements have the same uncertainty σ_i , the value of σ_i is allowed to vary over the course of the observation (i.e. with i). The basic idea is as follows:

The current best-estimate of the measurement uncertainty $\sigma_{\text{ref}, i-1}$ is used in (25) to estimate the value of TASTE at time t_i :

$$\text{TASTE}_i = \frac{2n_i (1 - \lambda_{\text{max}, i})}{\sigma_{\text{ref}, i-1}^2}. \tag{27}$$

Having eliminated any bad quality measurements, the fit may be assumed to be good and the resulting value of TASTE_i may be used to estimate σ_i :

$$\hat{\sigma}_i = \sigma_{\text{ref}, i-1} \sqrt{\frac{\text{TASTE}_i}{2n_i - 3}}. \tag{28}$$

This estimate is unbiased but, as the number of measurements is small, the variance of its error will be large [see 20, pp. 5–7]. To improve the accuracy of the value used in (27), whilst at the same time permitting the measurement uncertainty to vary over an observation (as new stars enter and leave the field-of-view), exponential smoothing is used to effectively implement a very simple low-pass filter, i.e.

$$\sigma_{\text{ref}, i} = \alpha \hat{\sigma}_i + (1 - \alpha) \sigma_{\text{ref}, i-1}, \quad i > 10, \tag{29}$$

where α is the smoothing factor.³¹ To initialize the smoother, whilst avoiding giving too much weight to the first few estimates, a moving average is used:

$$\sigma_{\text{ref}, i} = \frac{\hat{\sigma}_i + i \sigma_{\text{ref}, i-1}}{i + 1}, \quad i = 1, \dots, 10, \tag{30}$$

³¹The smoothing factor is nominally set to the value 0.1, but may be modified by means of the parameter `alpha` (see p. 29, Table 2).



where $\sigma_{\text{ref},0}$ is nominally set to $3''$, but may be modified by means of the parameter `measerr_init` (see p. 29, Table 2).

Calculation of the error covariance matrix

Following Shuster and Oh [21, p. 71], we define the Cartesian attitude covariance matrix as

$$P_{\theta\theta} = E(\delta\theta \delta\theta^T),$$

where $\delta\theta = (\delta\theta_1 \ \delta\theta_2 \ \delta\theta_3)^T$ and the $\delta\theta_i$ are the small-angle rotations about the body (ACA-frame) axes which take the true spacecraft attitude to its estimated attitude. Transforming the expression obtained in [21] from the BRF-frame to the ACA-frame we obtain:

$$P_{\theta\theta} = \sigma^2 A_{\text{align}} \left[\sum_{i=1}^n (I - \mathbf{u}'_i \mathbf{u}'_i{}^T) \right]^{-1} A_{\text{align}}^T, \quad (31)$$

where $A_{\text{align}} = A(\mathbf{q}_{\text{align}})$.³² (As in the section describing the estimation of the measurement uncertainty, it is assumed that any bad star measurements have already been eliminated.)

³²The expression for the covariance matrix in the BRF follows from eqs. (87) and (99) of [21, p. 75], where it is also noted that a singular matrix $\sum_{i=1}^n (I - \mathbf{u}'_i \mathbf{u}'_i{}^T)$ would signal a non-unique solution for the optimal quaternion $\hat{\mathbf{q}}$.



2.2 Gyro-based attitude reconstruction

2.2.1 Derivation of the model

Suppose we wish to estimate the (ACA-frame) attitude of the spacecraft, $\mathbf{q}_{\text{aca}}(t)$, at the time, t_k , of the k 'th gyro measurement. We start by choosing a reference attitude $\mathbf{q}_{0,k}$ such that, for t sufficiently close to t_k , the rotation between $\mathbf{q}_{0,k}$ and $\mathbf{q}_{\text{aca}}(t)$, i.e. the angular rotation represented by the quaternion

$$\mathbf{q}_r^{(k)}(t) \equiv \mathbf{q}_{0,k}^* \mathbf{q}_{\text{aca}}(t), \quad (32)$$

is small (* denotes the conjugate operation).³³

Suppose $\mathbf{q}_r^{(k)}(t)$ corresponds to the rotation $\theta^{(k)}(t)$ about the unit vector $\mathbf{e}^{(k)}(t)$. Dropping the explicit references to t and k , we may write

$$\mathbf{q}_r = \begin{bmatrix} e_x \sin \frac{\theta}{2} & e_y \sin \frac{\theta}{2} & e_z \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{bmatrix}^T, \quad (33)$$

where e_x , e_y and e_z are the components of \mathbf{e} with respect to the $\mathbf{q}_{0,k}$ reference frame. For small θ , we may approximate $\sin \frac{\theta}{2}$ by $\frac{\theta}{2}$ and $\cos \frac{\theta}{2}$ by 1, so that

$$\mathbf{q}_r = \begin{bmatrix} \frac{\theta_x}{2} & \frac{\theta_y}{2} & \frac{\theta_z}{2} & 1 \end{bmatrix}^T + O(\theta^2), \quad (34)$$

where $\theta_x = \theta e_x$, $\theta_y = \theta e_y$ and $\theta_z = \theta e_z$, and the kinematic equations of motion [e.g. 26, p. 512]

$$\frac{d\mathbf{q}_r}{dt} = \frac{1}{2} \Omega \mathbf{q}_r, \quad \Omega = \begin{bmatrix} 0 & \omega_w & -\omega_v & \omega_u \\ -\omega_w & 0 & \omega_u & \omega_v \\ \omega_v & -\omega_u & 0 & \omega_w \\ -\omega_u & -\omega_v & -\omega_w & 0 \end{bmatrix},$$

³³To achieve this, `calcGyroAttitude` sets $\mathbf{q}_{0,k} = \hat{\mathbf{q}}_{\text{str}}^{(m)}$, where $\hat{\mathbf{q}}_{\text{str}}^{(m)}$ is the estimate of the ACA-attitude based on the star tracker data at time $t_s^{(m)}$ and $m = \max\{i : t_s^{(i)} \leq t_k\}$, whenever it is found that the rotation represented by $\mathbf{q}_{0,k}^* \hat{\mathbf{q}}_{\text{str}}^{(m)}$ exceeds a certain threshold, which the User may specify by means of the input parameter `ref_thresh`. In addition, it is ensured that only good quality attitude measurements are used as the reference attitude (if the initial attitude measurement is of bad quality—which may indicate that it could not be estimated—then the on-board attitude estimate is used).



become:

$$\begin{aligned}
 \dot{\theta}_x &= \omega_u + \frac{1}{2}[\omega_w \theta_y - \omega_v \theta_z] + O(\theta^2), \\
 \dot{\theta}_y &= \omega_v + \frac{1}{2}[\omega_u \theta_z - \omega_w \theta_x] + O(\theta^2), \\
 \dot{\theta}_z &= \omega_w + \frac{1}{2}[\omega_v \theta_x - \omega_u \theta_y] + O(\theta^2),
 \end{aligned} \tag{35}$$

where ω_u , ω_v and ω_w are the components of the angular velocity vector in the body-fixed (i.e. ACA) reference frame and clearly $\theta = \sqrt{\theta_x^2 + \theta_y^2 + \theta_z^2}$.

The determination of the body rates, ω_u , ω_v and ω_w , from the rates, ω_1 , ω_2 , ω_3 and ω_4 , about each of the four gyro axes is an overdetermined problem. We may either disregard the measurements from one of the gyros, using them instead as a check on the quality of the measurements from the other three, or, as is currently performed in the pointing reconstruction software, find the least-squares solution.³⁴ That is,

$$\begin{bmatrix} \omega_u \\ \omega_v \\ \omega_w \end{bmatrix} = G^+ \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \\ \omega_4 \end{bmatrix}, \tag{36}$$

where $G^+ = (G^T G)^{-1} G^T$ is the Moore-Penrose pseudoinverse of the gyro alignment matrix G .³⁵

Substituting (36) in (35) and neglecting terms of degree one in θ ,³⁶ we

³⁴It is argued in HCSS-19121 that it might be preferable to follow the method that was used on-board and disregard the measurements from one of the four gyros.

³⁵For the nominal (pre-launch) alignments, we have $G = \frac{1}{\sqrt{3}} \begin{pmatrix} -1 & -1 & 1 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \\ -1 & -1 & -1 \end{pmatrix}$ and $G^+ = \frac{\sqrt{3}}{4} \begin{pmatrix} -1 & 1 & 1 & -1 \\ -1 & -1 & -1 & -1 \\ 1 & 1 & -1 & -1 \end{pmatrix}$ [1, pp. 599–600].

³⁶Writing $\underline{\theta} = \theta \mathbf{e}$ and $\underline{\omega} = [\omega_u \ \omega_v \ \omega_w]^T$, (35) become

$$\underline{\dot{\theta}} = \underline{\omega} + \frac{1}{2} \underline{\theta} \times \underline{\omega} + O(\theta^2).$$



obtain

$$\begin{bmatrix} \dot{\theta}_x \\ \dot{\theta}_y \\ \dot{\theta}_z \end{bmatrix} = G^+ \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \\ \omega_4 \end{bmatrix}. \quad (37)$$

Now, the Inertial Reference Unit (GYR) provides measurements of the integrated angular rotations, ϕ_1 , ϕ_2 , ϕ_3 and ϕ_4 , of the spacecraft about each of the four gyro axes (from some unknown time). If we assume that their derivatives are related to the true rates about these axes by

$$\dot{\phi}_i = k_i \omega_i + b_i, \quad i = 1, \dots, 4,$$

where k_i are the (known) scale factors and b_i are the (unknown) drift rates,³⁷ then (37) becomes

$$\begin{aligned} \dot{\theta}_x &= \dot{\psi}_x + b_x, \\ \dot{\theta}_y &= \dot{\psi}_y + b_y, \\ \dot{\theta}_z &= \dot{\psi}_z + b_z, \end{aligned} \quad (38)$$

where

$$\begin{bmatrix} \psi_x \\ \psi_y \\ \psi_z \end{bmatrix} \equiv G^+ \begin{bmatrix} \phi_1/k_1 \\ \phi_2/k_2 \\ \phi_3/k_3 \\ \phi_4/k_4 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} b_x \\ b_y \\ b_z \end{bmatrix} = -G^+ \begin{bmatrix} b_1/k_1 \\ b_2/k_2 \\ b_3/k_3 \\ b_4/k_4 \end{bmatrix}. \quad (39)$$

Next, we assume that a value T be chosen such that the variation in the gyro drift rates (and hence in b_x , b_y and b_z) is negligible within the interval $I_k \equiv [t_k - T/2, t_k + T/2]$.³⁸ For $t \in I_k$, equations (38) can then be integrated

So neglecting the linear term in θ introduces an error in $\dot{\theta}$ of magnitude $\dot{\theta}_{\text{err}}$, where

$$\frac{\dot{\theta}_{\text{err}}}{\dot{\theta}} \approx \frac{|\underline{\theta} \times \underline{\omega}|}{2|\underline{\omega}|} \leq \frac{\theta}{2}.$$

For example, for $\theta = 1^\circ$, $\dot{\theta}_{\text{err}}/\dot{\theta} < 0.01$.

³⁷The function `calcGyroAttitude` reads the values of k_i from the ACMS product metadata.

³⁸ The appropriateness of using this ‘constant drift rate’ model remains to be investigated. (Alternative models, such as one in which the gyro drift rates vary linearly with the time, could be easily implemented.) The window length, T , is specified by the input parameter `wind_len`, the default value for which is currently 400 s; see [9, p. 7] and [24, p. 68].



to give:

$$\begin{aligned}
 \theta_x(t) &= \psi_x(t) + b_x(t - t_k) + c_x, \\
 \theta_y(t) &= \psi_y(t) + b_y(t - t_k) + c_y, \\
 \theta_z(t) &= \psi_z(t) + b_z(t - t_k) + c_z,
 \end{aligned}
 \tag{40}$$

where c_x, c_y, c_z are constants of integration.

Equations (40) constitute the model used to relate the small rotations, θ_x, θ_y and θ_z , from the reference attitude $\mathbf{q}_{0,k}$ to three angles, ψ_x, ψ_y and ψ_z , derived from the gyro output.

2.2.2 Estimation of the model parameters

To find (estimate) the parameters in (40) we make use of the improved attitude estimates constructed from the star tracker data (see Section 2.1). Measurements of ψ_x, ψ_y and ψ_z are combined with measurements of θ_x, θ_y and θ_z and a least-squares fit is performed.

Let t_b and t_e be the start and end times of the observation and suppose that within the interval $I_k \cap [t_b, t_e]$ we have n_g measurements of the angles ψ_x, ψ_y and ψ_z at times $t_g^{(1)}, \dots, t_g^{(n_g)}$, i.e.

$$\{(\psi_x^{(i)}, \psi_y^{(i)}, \psi_z^{(i)}) : i = 1, \dots, n_g\},$$

and n_s good quality measurements of the angles θ_x, θ_y and θ_z at times $t_s^{(1)}, \dots, t_s^{(n_s)}$,³⁹ i.e.

$$\{(\theta_x^{(i)}, \theta_y^{(i)}, \theta_z^{(i)}) : i = 1, \dots, n_s\}.$$

Since $t_k \in I_k$, we may without loss of generality suppose that

$$t_k = t_g^{(l)}, \tag{41}$$

³⁹ To allow the timing of the star tracker measurements to be synchronized with that of the gyro measurements, an offset set is first added, i.e. $t_s^{(i)} = \tilde{t}_s^{(i)} + \text{toff_star}$, where $\tilde{t}_s^{(i)}$ is the time of the measurement created by `calcStrAttitude` and `toff_star` is an input parameter which may be specified by the User. It is noted that when assigning the time to each attitude measurement the function `calcStrAttitude` subtracts 0.805 s from the record time associated with the particular set of star unit vectors. Therefore, the total offset is in fact `toff_star` - 0.805 s. The measurements $\theta_x^{(i)}, \theta_y^{(i)}$ and $\theta_z^{(i)}$ are considered to be of good quality whenever the p_{taste} associated with the improved attitude estimate $\hat{\mathbf{q}}_{\text{str}}^{(i)}$, see (26), is greater than the value of the input parameter `prob_thresh`.



for some $1 \leq l \leq n_g$.

To obtain the measurements, $\theta_x^{(i)}$, $\theta_y^{(i)}$, $\theta_z^{(i)}$, from the reference attitude $\mathbf{q}_{0,k}$ and the improved attitude estimate $\hat{\mathbf{q}}_{\text{str}}^{(i)}$ we simply use (32) and (34), i.e.

$$\begin{bmatrix} \frac{\theta_x^{(i)}}{2} & \frac{\theta_y^{(i)}}{2} & \frac{\theta_z^{(i)}}{2} & 1 \end{bmatrix}^T = \mathbf{q}_{0,k}^* \hat{\mathbf{q}}_{\text{str}}^{(i)}. \quad (42)$$

In doing this we are approximating $\sin \frac{\theta^{(i)}}{2}$ by $\frac{\theta^{(i)}}{2}$ (see p. 17) and so introducing errors of up to $2 \left| \sin \frac{\theta^{(i)}}{2} - \frac{\theta^{(i)}}{2} \right| = \frac{(\theta^{(i)})^3}{24} + O([\theta^{(i)}]^5)$, i.e. approximately $0.006''$ ($\theta^{(i)} = 0.5^\circ$), $0.05''$ ($\theta^{(i)} = 1^\circ$) and $0.4''$ ($\theta^{(i)} = 2^\circ$). To ensure that these systematic errors do not corrupt the results, any measurement for which the rotation, $\theta^{(i)}$, from the reference attitude is greater than the parameter `rot_limit` is not used in the fitting.⁴⁰

Typically there are four times as many measurements of ψ_x , ψ_y and ψ_z as of θ_x , θ_y and θ_z , i.e. $n_g \approx 4n_s$, and the times $\{t_g^{(i)} : i = 1, \dots, n_g\}$ and $\{t_s^{(i)} : i = 1, \dots, n_s\}$ do not ‘coincide’. Moreover, the noise on the star tracker measurements, and hence on $\theta_x^{(i)}$, $\theta_y^{(i)}$ and $\theta_z^{(i)}$, is far greater than that arising from the gyro measurements. It follows that, to combine the two sets of measurements, it is appropriate to interpolate the measurements $\psi_x^{(i)}$, $\psi_y^{(i)}$, $\psi_z^{(i)}$ to the times $\{t_s^{(i)} : i = 1, \dots, n_s\}$ of the star tracker measurements. For this, the software uses linear interpolation, i.e.

$$\tilde{\psi}_x^{(i)} = \frac{\psi_x^{(j+1)} [t_s^{(i)} - t_g^{(j)}] + \psi_x^{(j)} [t_g^{(j+1)} - t_s^{(i)}]}{t_g^{(j+1)} - t_g^{(j)}}, \quad t_g^{(j)} \leq t_s^{(i)} < t_g^{(j+1)},$$

and similarly for $\tilde{\psi}_y^{(i)}$ and $\tilde{\psi}_z^{(i)}$.

Using this set of $3n_s$ measurements with the model (40), we obtain the

⁴⁰This is achieved by assigning the measurement a zero weight when solving (45) and incrementing the number, n_d , of discarded measurements.



following linear regression equations:

$$\begin{bmatrix} \theta_x^{(1)} - \tilde{\psi}_x^{(1)} \\ \theta_y^{(1)} - \tilde{\psi}_y^{(1)} \\ \theta_z^{(1)} - \tilde{\psi}_z^{(1)} \\ \dots \\ \theta_x^{(n_s)} - \tilde{\psi}_x^{(n_s)} \\ \theta_y^{(n_s)} - \tilde{\psi}_y^{(n_s)} \\ \theta_z^{(n_s)} - \tilde{\psi}_z^{(n_s)} \end{bmatrix} = \begin{bmatrix} t_s^{(1)} - t_k & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & t_s^{(1)} - t_k & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & t_s^{(1)} - t_k & 1 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ t_s^{(n_s)} - t_k & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & t_s^{(n_s)} - t_k & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & t_s^{(n_s)} - t_k & 1 \end{bmatrix} \begin{bmatrix} b_x \\ c_x \\ b_y \\ c_y \\ b_z \\ c_z \end{bmatrix} + \begin{bmatrix} \epsilon_x^{(1)} \\ \epsilon_y^{(1)} \\ \epsilon_z^{(1)} \\ \dots \\ \epsilon_x^{(n_s)} \\ \epsilon_y^{(n_s)} \\ \epsilon_z^{(n_s)} \end{bmatrix},$$

which we write, more concisely, as

$$y = X \beta + \epsilon. \tag{43}$$

The error term, ϵ , is assumed to have a (conditional) mean of zero and (conditional) variance given by the block diagonal matrix:

$$\Omega = \text{diag}(P_{\theta\theta}^{(1)}, \dots, P_{\theta\theta}^{(n_s)}),$$

where $P_{\theta\theta}^{(i)}$ is the Cartesian attitude covariance matrix of measurement i , computed according to (31).

To find the best linear unbiased estimate of β (i.e. of the parameters $b_x, c_x, \dots, b_z, c_z$) we minimize, with respect to β , the cost function $J(\beta) = \epsilon^T \Omega^{-1} \epsilon$, i.e. we use the generalized least-squares method. Since Ω is symmetric and positive definite, the Cholesky decomposition may be used to write $\Omega = L L^T$, where L is lower triangular. Setting $y^* = L^{-1} y$, $X^* = L^{-1} X$ and $\epsilon^* = L^{-1} \epsilon$, it follows that

$$\begin{aligned} J(\beta) &= (y - X \beta)^T \Omega^{-1} (y - X \beta) \\ &= (L y^* - L X^* \beta)^T (L L^T)^{-1} (L y^* - L X^* \beta) \\ &= (y^* - X^* \beta)^T (y^* - X^* \beta) \\ &= (\epsilon^*)^T \epsilon^*, \end{aligned}$$

and the problem becomes one of ordinary least-squares. Our estimate of β may thus be found by solving the normal equations

$$[(X^*)^T X^*] \hat{\beta} = (X^*)^T y^*. \tag{44}$$



Instead of solving (44), the software assumes that the off-diagonal elements of $P_{\theta\theta}^{(i)}$ are negligible.⁴¹ Ω and L then become diagonal and equations (44) decouple to give:

$$\begin{aligned} (X_x^T X_x) \begin{bmatrix} \hat{b}_x \\ \hat{c}_x \end{bmatrix} &= X_x^T y_x, \\ (X_y^T X_y) \begin{bmatrix} \hat{b}_y \\ \hat{c}_y \end{bmatrix} &= X_y^T y_y, \\ (X_z^T X_z) \begin{bmatrix} \hat{b}_z \\ \hat{c}_z \end{bmatrix} &= X_z^T y_z, \end{aligned} \quad (45)$$

where

$$X_x = \begin{bmatrix} (t_s^{(1)} - t_k)/\sigma_x^{(1)} & 1/\sigma_x^{(1)} \\ (t_s^{(2)} - t_k)/\sigma_x^{(2)} & 1/\sigma_x^{(2)} \\ \dots & \dots \\ (t_s^{(n_s)} - t_k)/\sigma_x^{(n_s)} & 1/\sigma_x^{(n_s)} \end{bmatrix}, \quad y_x = \begin{bmatrix} (\theta_x^{(1)} - \tilde{\psi}_x^{(1)})/\sigma_x^{(1)} \\ (\theta_x^{(2)} - \tilde{\psi}_x^{(2)})/\sigma_x^{(2)} \\ \dots \\ (\theta_x^{(n_s)} - \tilde{\psi}_x^{(n_s)})/\sigma_x^{(n_s)} \end{bmatrix},$$

$\sigma_x^{(i)} = \sqrt{P_{\theta\theta}^{(i)}[1, 1]}$ and similiary for X_y, y_y, σ_y and X_z, y_z, σ_z .⁴²

2.2.3 Calculation of the goodness-of-fit

Solving (45) is equivalent to minimizing the cost functions:

$$J_r(b_r, c_r) = \left(y_r - X_r \begin{bmatrix} b_r \\ c_r \end{bmatrix} \right)^T \left(y_r - X_r \begin{bmatrix} b_r \\ c_r \end{bmatrix} \right), \quad r \in \{x, y, z\}. \quad (46)$$

If our model (40) accurately reflects reality and if the measurement errors are as expected, i.e. if the (normalized) errors contributing to the right-hand sides of (46) come from the standard normal distribution (see Appendix C), then it follows that each $J_r(\hat{b}_r, \hat{c}_r)$ will come from a chi-squared distribution with $n_s - n_d - 2$ degrees of freedom [16, p. 654].⁴³ That is, the probability of obtaining a value of the minimized cost function which exceeds $\chi_r^2(\hat{b}_r, \hat{c}_r)$ is

$$p_r \equiv \Pr\{\min\{J_r\} > J_r(\hat{b}_r, \hat{c}_r)\} = 1 - P\left(\frac{n_s - n_d - 2}{2}, \frac{\chi_r^2(\hat{b}_r, \hat{c}_r)}{2}\right), \quad (47)$$

⁴¹The impact of making this simplification is investigated in Appendix B.

⁴² $P_{\theta\theta}^{(i)}[j, k]$ indicates the jk^{th} component of $P_{\theta\theta}^{(i)}$.

⁴³Remember that n_d is the number of discarded measurements (see footnote 40).



where $r \in \{x, y, z\}$ and P once again denotes the lower (regularized) incomplete gamma function.

The three p -values, p_x , p_y and p_z , may be interpreted as the likelihoods, given the data, associated with the each of the three sets of estimated parameters $\{\hat{b}_x, \hat{c}_x\}$, $\{\hat{b}_y, \hat{c}_y\}$ and $\{\hat{b}_z, \hat{c}_z\}$. However, what we require is the likelihood that all three sets of parameters are simultaneously correct. To achieve this we use Fisher's method [11, pp. 103–105] and calculate the value of the test statistic

$$T = -2 \ln(p_x p_y p_z). \quad (48)$$

It may be easily shown that T has a chi-squared distribution with 6 degrees-of-freedom, so that the combined p -value is

$$p_{\text{comb'd}} = 1 - P\left(3, \frac{T}{2}\right). \quad (49)$$

2.2.4 Computation of the new attitude estimate and its uncertainty

Assuming the quality of fitting, as indicated by $p_{\text{comb'd}}$, is good, (40) and (41) may be used to obtain estimates of the small-angle rotations at time t_k , i.e.

$$\hat{\theta}_r(t_k) = \psi_r^{(l)} + \hat{c}_r, \quad r \in \{x, y, z\}. \quad (50)$$

Then from (34),

$$\hat{\mathbf{q}}_r^{(k)}(t_k) = \begin{bmatrix} \frac{\hat{\theta}_x(t_k)}{2} & \frac{\hat{\theta}_y(t_k)}{2} & \frac{\hat{\theta}_z(t_k)}{2} & 1 \end{bmatrix}^T, \quad (51)$$

and, finally, from (32),

$$\hat{\mathbf{q}}_{\text{aca}}(t_k) = \mathbf{q}_{0,k} \hat{\mathbf{q}}_r^{(k)}(t_k). \quad (52)$$

Let $\tilde{\theta}_{x,k}$, $\tilde{\theta}_{y,k}$ and $\tilde{\theta}_{z,k}$ be the errors in the small angle rotations at time t_k . In passing from (44) to (45) we have ignored the coupling between these errors and so we are now unable to compute $E[\tilde{\theta}_{x,k} \tilde{\theta}_{y,k}]$, $E[\tilde{\theta}_{x,k} \tilde{\theta}_{z,k}]$ and $E[\tilde{\theta}_{y,k} \tilde{\theta}_{z,k}]$. However, the expected variances of the errors in c_x , c_y and c_z are still available and may be readily obtained from the diagonal elements of the Cartesian attitude covariance matrix contain, i.e.

$$E[\tilde{\theta}_{r,k}^2] = E(\tilde{c}_r^2) = (X_r^T X_r)^{-1}[2, 2], \quad r \in \{x, y, z\}. \quad (53)$$



These elements may then be used to create confidence regions in which the true attitude is believed to lie. In Appendix B, an elliptical confidence region for the direction of the spacecraft x -axis is calculated, both using the full, coupled system of equations and the decoupled system. For the particular dataset studied, the difference was found to be very small.



3 How to use the software

To obtain a pointing product containing the attitudes reconstructed by the gyro-based method, use `calcAttitude`. This task may be either selected and run from the task view or called from the command line:

```
newpp = calcAttitude (oldpp, acmsProduct, tcHistoryProduct,  
                    [prob_thresh = ..., prob_frac = ...,  
                    star_set = ..., measerr_init = ..., alpha = ...,  
                    ref_thresh = ..., wind_len = ..., rot_limit = ...,  
                    toff_star = ..., excl_gyro = ...])
```

By calling first `calcStrAttitude` and then `calcGyroAttitude`, the task `calcAttitude` takes the information contained in the ACMS telemetry product, `acmsProduct`, and uses the method described in Section 2 to estimate the ACA-frame attitude at a sequence of times spanning the period of the original pointing product, `oldpp`.⁴⁴ It then creates a new copy, `newpp`, of the pointing product and augments this with the newly estimated attitudes and the associated quality and accuracy information. To do this, the task loops through the lines of the pointing product and replaces each of the attitude quaternions in the column labeled `filterQuat` with the newly-estimated attitude, $\hat{\mathbf{q}}_{aca}(t_k)$, associated with the closest matching time, t_k . The related quality and accuracy information are written to the columns labelled `gyroAttProbX`, `gyroAttProbY`, `gyroAttProbZ` and `gyroAttSigmaX`, `gyroAttSigmaY`, `gyroAttSigmaZ` respectively. To indicate that this has been performed, the flag in the column labelled `filterQuatFlag` is set equal to 0. If, for some time t_k , the gyro-based method was unable to estimate the attitude, then the quaternion in the column labeled `filterQuat` and the flag in the column labelled `filterQuatFlag` are left unchanged. Irrespective of whether the values in column `filterQuat` are overwritten, the old attitude quaternion—which is that computed by the on-board filter (with or without the simple focal length correction)—is copied to the column labelled `simpleCorrFilterQuat`.

⁴⁴As explained in Section 2.2, the gyro-based method reconstructs the attitudes at the times associated with the gyro measurements. Although these times are very similar to the OBT times in the pointing product, they do not match exactly (see also Section 4, HCSS-19201). The TC history product, `tcHistoryProduct`, is also required as it contains the spacecraft velocity vector which is needed in the aberration correction.



A summary of the columns in the pointing product associated with the gyro-based attitude reconstruction is given below in Table 1; for a complete description of all the columns, see [17]. In addition, `calcStrAttitude` accepts a number of optional input parameters; these are described in Table 2.

Column	Field	Units	Description
1	obt	μs	On-board time.
3	filterQuat	None	Reconstructed ACA-frame attitude quaternion. Iff <code>filterQuatFlag = 0</code> , then <code>filterQuat</code> contains $\hat{\mathbf{q}}_{aca}(t_k)$, where t_k is the time of the closest gyro measurement.
19	gyroAttProbX	None	Quality associated with x -axis fit, i.e. $\Pr\{\min\{\chi_x^2\} > \chi_x^2(\hat{b}_x, \hat{c}_x)\}$. ⁴⁵
20	gyroAttProbY	None	Quality associated with y -axis fit, i.e. $\Pr\{\min\{\chi_y^2\} > \chi_y^2(\hat{b}_y, \hat{c}_y)\}$.
21	gyroAttProbZ	None	Quality associated with z -axis fit, i.e. $\Pr\{\min\{\chi_z^2\} > \chi_z^2(\hat{b}_z, \hat{c}_z)\}$.
22	gyroAttSigmaX	arcsec.	Standard deviation of error in gyro-based reconstructed attitude about ACA-frame x -axis, i.e. $\sqrt{E[\tilde{\theta}_{x,k}^2]}$. ⁴⁶
23	gyroAttSigmaY	arcsec.	Standard deviation of error in gyro-based reconstructed attitude about ACA-frame y -axis, i.e. $\sqrt{E[\tilde{\theta}_{y,k}^2]}$.
24	gyroAttSigmaZ	arcsec.	Standard deviation of error in gyro-based reconstructed attitude about ACA-frame z -axis, i.e. $\sqrt{E[\tilde{\theta}_{z,k}^2]}$.
25	filterQuatFlag	None	Flag indicating nature of attitude in <code>filterQuat</code> field. (0 = gyro-based method, 1 = simple focal length correction, 2 = on-board filter)

Table 1: Relevant columns of pointing product

⁴⁵See p. 23, eq. (47).

⁴⁶See p. 24, eq. (53).



Parameter	Class	Description
oldpp	PointingProduct	Original pointing product, i.e. <code>obs.auxiliary.pointing</code> where <code>obs</code> is an observation context.
acmsProduct	AcmsTelemetryProduct	ACMS telemetry product, i.e. <code>obs.auxiliary.acms</code> where <code>obs</code> is an observation context.
tcHistoryProduct	TeleCommandHistProduct	TC history product, i.e. <code>obs.auxiliary.teleCommHistory</code> where <code>obs</code> is an observation context.
newpp	PointingProduct	New pointing product.
prob_thresh	Double	Probability threshold used for deciding when fit to determine star tracker attitude is bad (optional, def. value = 1.0×10^{-4}); see p. 14.
prob_frac ⁴⁷	Double	Probability factor used in ‘bad star’ test (optional, default value = 100.0); see p. 15
star_set	Int	Controls which stars are used to determine attitude: 0 = all, 1 = first nine, 2 = last nine (optional, default value = 0)
back_prop	Int	Controls which stars are back-propagated prior to amendment of distortion correction: 0 = all, 1 = first nine, 2 = last nine (optional, def. val. = 2); see foot. 16, p. 8.
measerr_init	Double	Initial estimate of star vector measurement error $\sigma_{ref,0}$ (optional, default value = 3.0''); see p. 16.
alpha	Double	Smoothing factor α for low-pass filtering of estimated measurement error (optional, default value = 0.1); see footnote 31, p. 15.
ref_thresh	Double	The reference attitude is updated whenever it is found to differ by more than <code>ref_thresh</code> (arcsecs.) from the ‘current’ star tracker attitude measurement (optional, default value = 100.0); see p. 17, footnote 33.

⁴⁷The `calcStrAttitude` header gives this parameter’s name as `prob_fac`, which would have been a more appropriate name!



wind_len	Double	The nominal length (in seconds) of the interval used to estimate the parameters in the least-squares fits (optional, default value = 400.0); see p. 19, footnote 38.
rot_limit	Double	Attitude measurements which differ by more than rot_limit (deg.) from the reference attitude are excluded from the fitting (optional, default value = 0.5); see p. 21.
toff_star	Double	Time offset (in seconds) used to synchronize the star tracker attitude measurements with the gyro measurements (optional, default value = 0.189); see p. 20, footnote 39. ⁴⁸
excl_gyro ⁴⁹	Int	Number of the gyro from which measurements are to be excluded. If set equal to zero, or any integer other than {1,2,3,4}, then measurements from all four gyros are used (optional, def. value = 0).
debug ⁵⁰	Int	Set equal to 1 for additional output (optional, default value = 0).
strAttitude	TableDataset	Contains corrected attitude measurements (from star tracker); see Table 3.
status	Int	Return status from calcStrAttitude (0 = success, otherwise error). See function header for further details.
gyroAttitude	TableDataset	Contains reconstructed attitude; see Table 4.

Table 2: Input and output parameters

⁴⁸The default value for `toff_star` is based upon the value which was estimated in [13] for being optimal for the now-defunct parameter `toff_gyro`; see HCSS-19454 for further details.

⁴⁹The `excl_gyro` option is no longer available in HIPE 15.

⁵⁰The `debug` option is no longer available in HIPE 15.



Before using the attitudes computed by the gyro-based method, it is advised that the User first checks the quality of the results. For example, a plot for observation 1342197884 of the probabilities found in columns 19–21 (Figure 1) indicates that the method was successful and that the reconstructed attitudes, as shown by the blue curves in Figures 2–4, are likely to be of good quality.⁵¹ (The combined probability, $p_{\text{comb'd}}$, when plotted is indistinguishable from the value 1.) The $1\text{-}\sigma$ uncertainties about the spacecraft axes are shown in Figure 5.

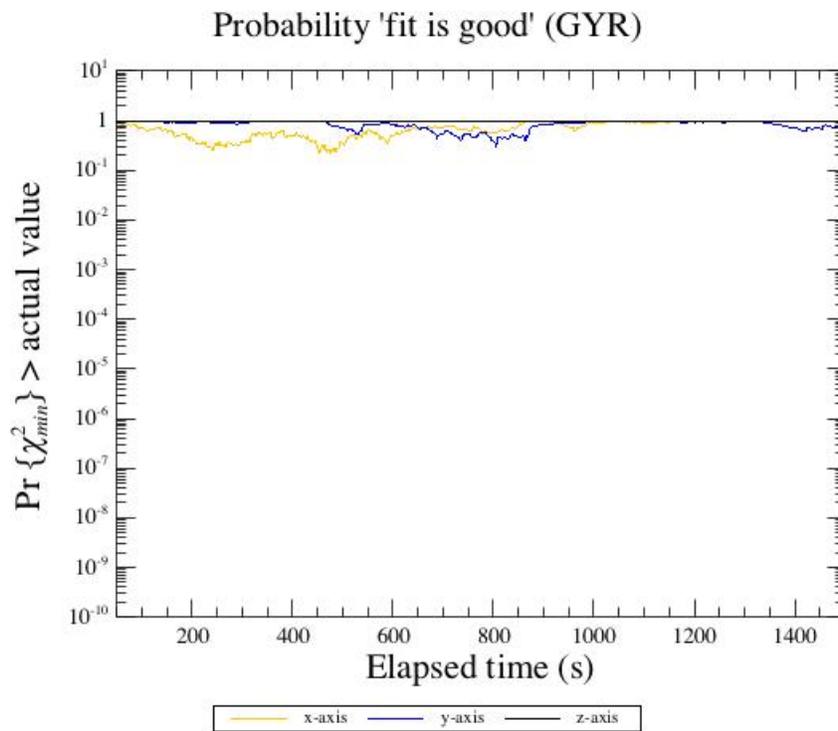


Figure 1: Quality plot (each axis) – 1342197884

⁵¹For this particular observation—in fact for much of OD 389—the attitudes from the gyro-based method are in closer agreement with those from the on-board filter than with the ‘simple-corrected’ attitudes. It is known that the ‘simple-corrected’ attitudes are likely to be inaccurate whenever interlacing mode was active; see HCSS-19870.

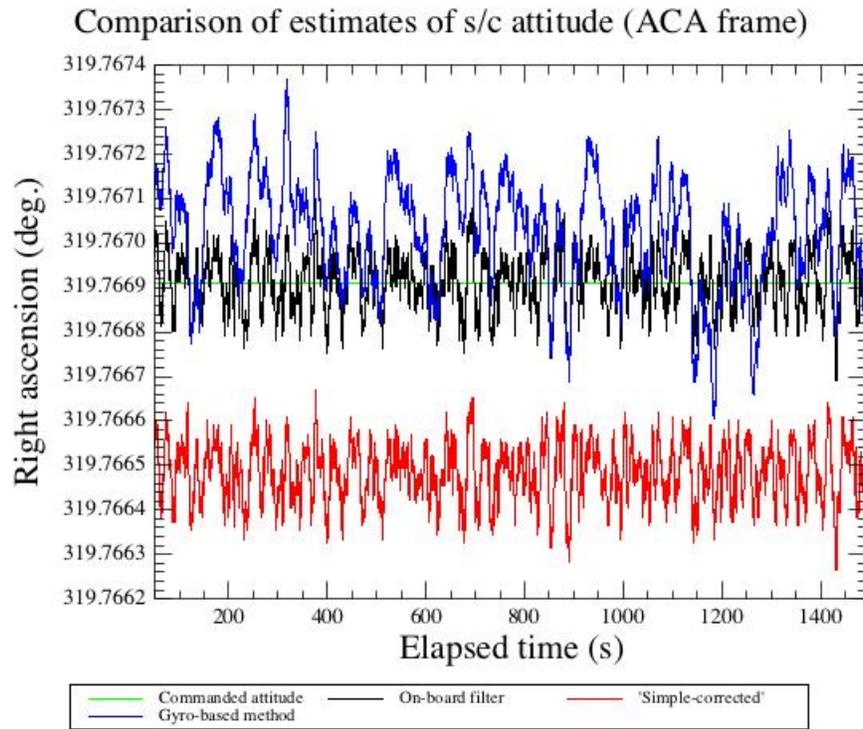


Figure 2: Right ascension – 1342197884

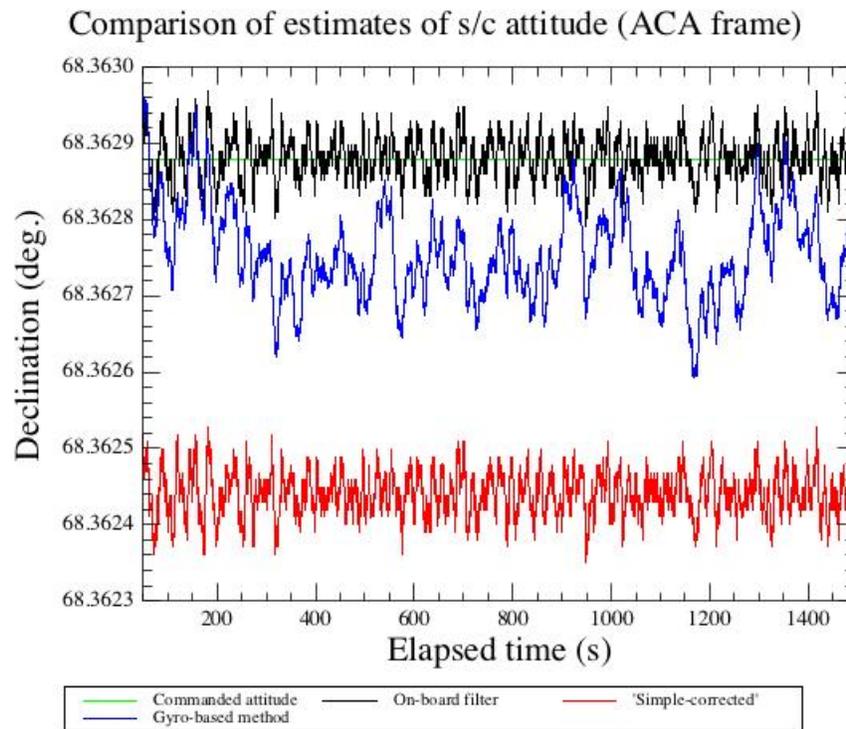


Figure 3: Declination – 1342197884

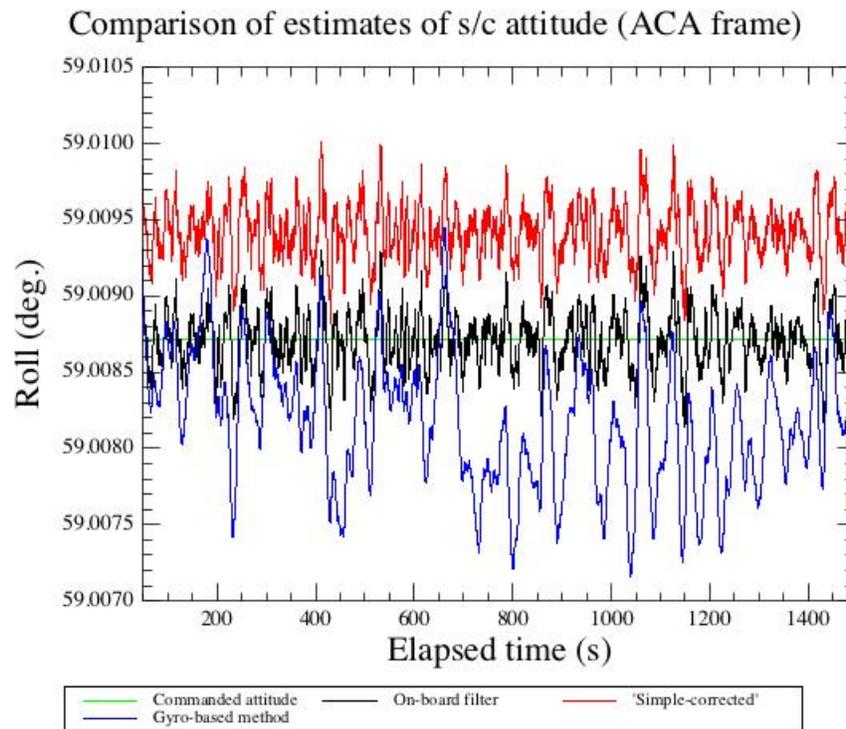


Figure 4: Roll – 1342197884

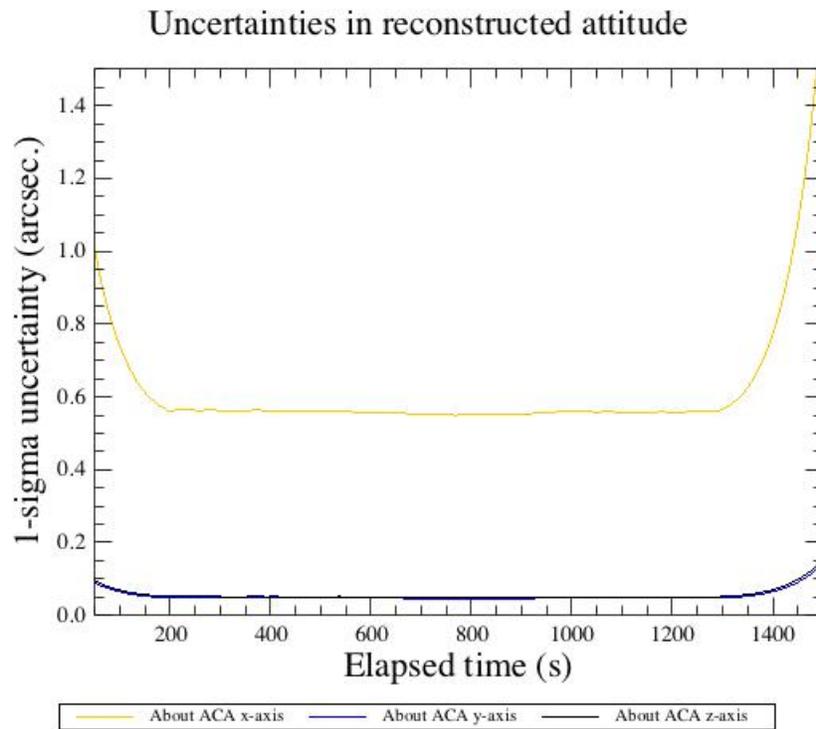


Figure 5: 1- σ uncertainties – 1342197884



Prior to the resolution of HCSS-19267 in August 2015, there appeared to be a problem with the pointing reconstruction for certain observations such as 1342227764 (see HCSS-19398). Figure 6 shows that, apart from during a short ~ 1000 s interval when interlacing mode was inactive, the quality appears to be very poor. For this observation (as, unfortunately, for

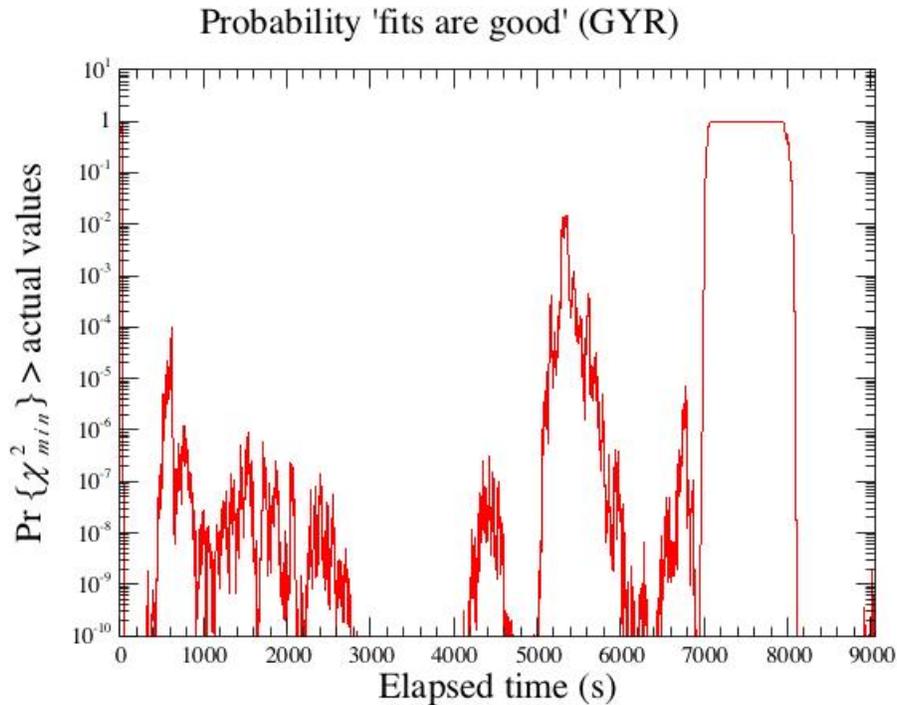


Figure 6: Quality plot (prior to build 2689) – 1342227764

all observations prior to OD 866)⁵² a maximum of nine stars are available in the `AcmsDtmStr` dataset and it appears that when interlacing mode was active the stars that are present are not the best nine. The problem was therefore that we were assuming the (mean) star vector measurement error to be $2.9''$ whereas, for much of this observation, the error was actually far worse (see Figure 7). Now that the software has been modified to estimate (and then use) the measurement error/uncertainty, the quality indicator no longer indicates there to be a problem (Figure 8). The presence of the lower

⁵²See p. 5, footnote 7.

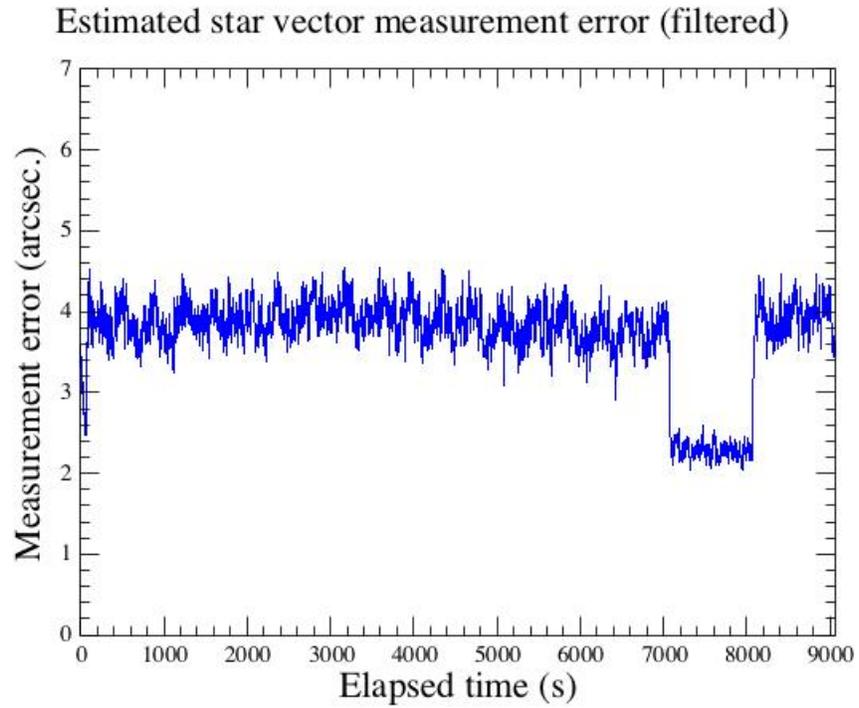


Figure 7: Estimated star vector measurement error – 1342227764

accuracy measurements is instead seen in the uncertainties associated with the reconstructed attitude (Figure 9).

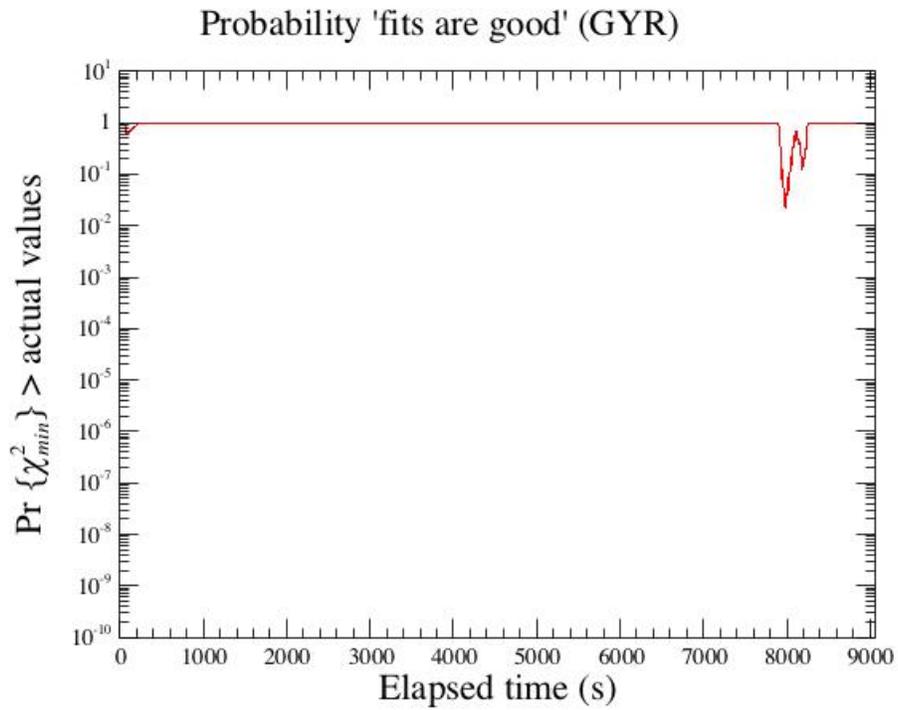


Figure 8: Quality plot (from build 2689 onwards) – 1342227764

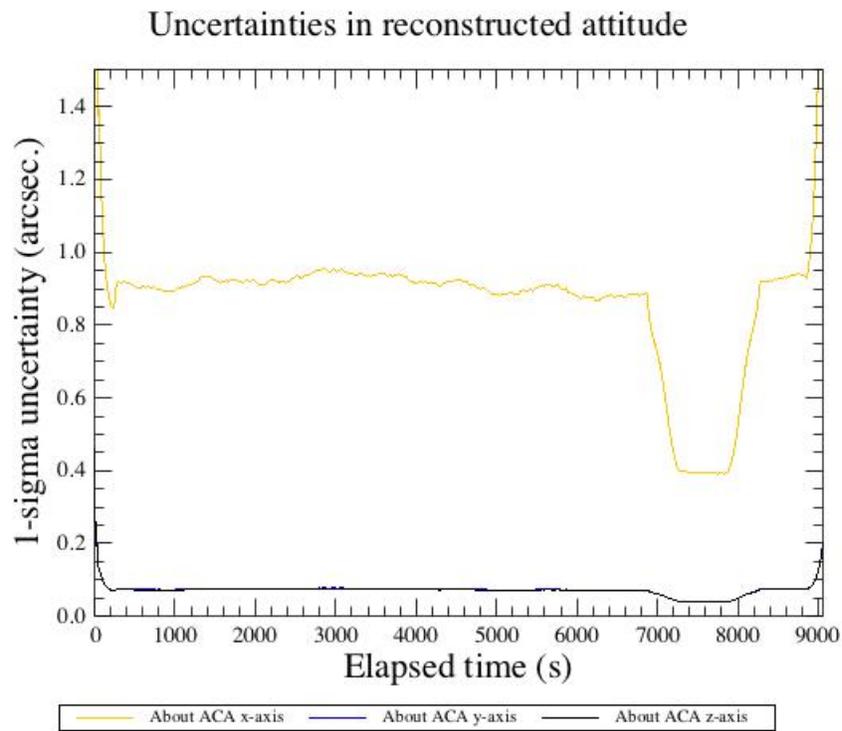


Figure 9: Quality plot (from build 2689 onwards) – 1342227764



To investigate cases such as this, additional output from the attitude reconstruction may sometimes be useful. This may be obtained by calling the functions `calcStrAttitude` and `calcGyroAttitude` directly and setting the debug flags,⁵³ i.e.

```
(strAttitude, status) = calcStrAttitude (oldpp, acmsProduct,
                                         tcHistoryProduct,
                                         [prob_thresh = ..., prob_frac = ...,
                                          star_set = ..., back_prop = ...,
                                          measerr_init = ..., alpha = ...,
                                          debug = 1])
```

```
if (status == 0):
```

```
    gyroAttitude = calcGyroAttitude (oldpp, acmsProduct, strAttitude,
                                      [ref_thresh = ..., wind_len = ...,
                                       rot_limit = ..., toff_star = ...,
                                       prob_thresh = ..., excl_gyro = ...,
                                       debug = 1])
```

The first call produces a table, `strAttitude`, containing the corrected attitude measurements made by the (operational) star tracker. These are then passed to the function `calcGyroAttitude`, which produces a table `gyroAttitude` containing the reconstructed attitude. Again, a full description of the various input and output parameters for these two calls is given in Table 2.⁵⁴

As an example, the quality of the attitude measurements produced by `calcStrAttitude` may be examined by means of the probabilities found in column 11 (`prob_taste`) and column 38 (`prob_taste_old`) of `strAttitude` (see Table 3). The first corresponds to the measurements produced after applying the new distortion correction and the second to measurements created using the uncorrected star vectors. Comparing plots of these quantities for observation 1342197884 the improvement achieved is clear (Figures 10 and 11).

⁵³As of HIPE 15.0, this is no longer possible (see p. 4 of the Introduction).

⁵⁴Unlike `calcAttitude`, this does not produce a new pointing product, but it has the advantage that the time stamps of the quaternions are slightly more accurate (see Section 4, HCSS-19201).

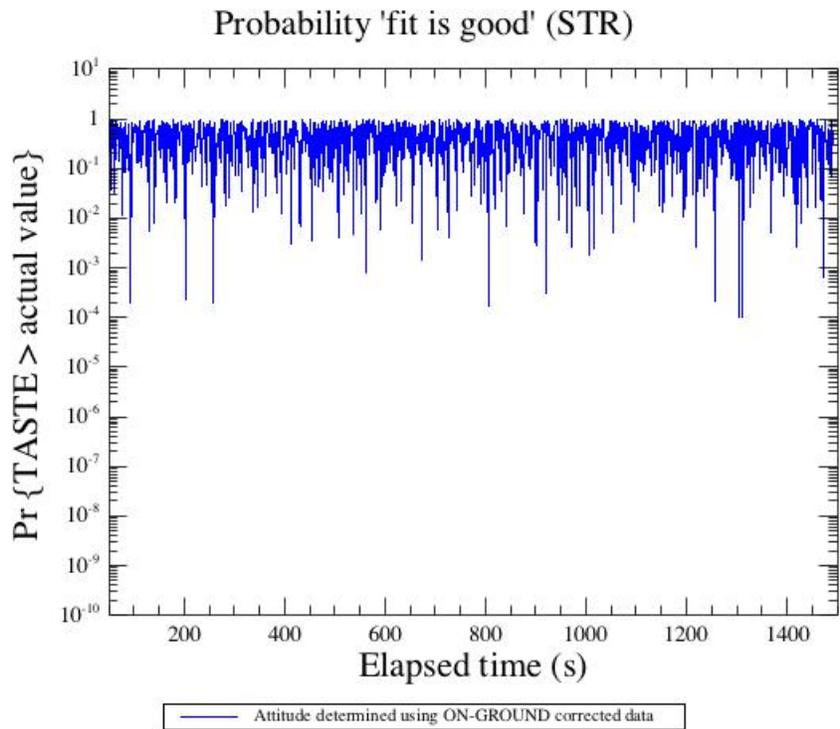


Figure 10: Quality of attitude measurements – 1342197884

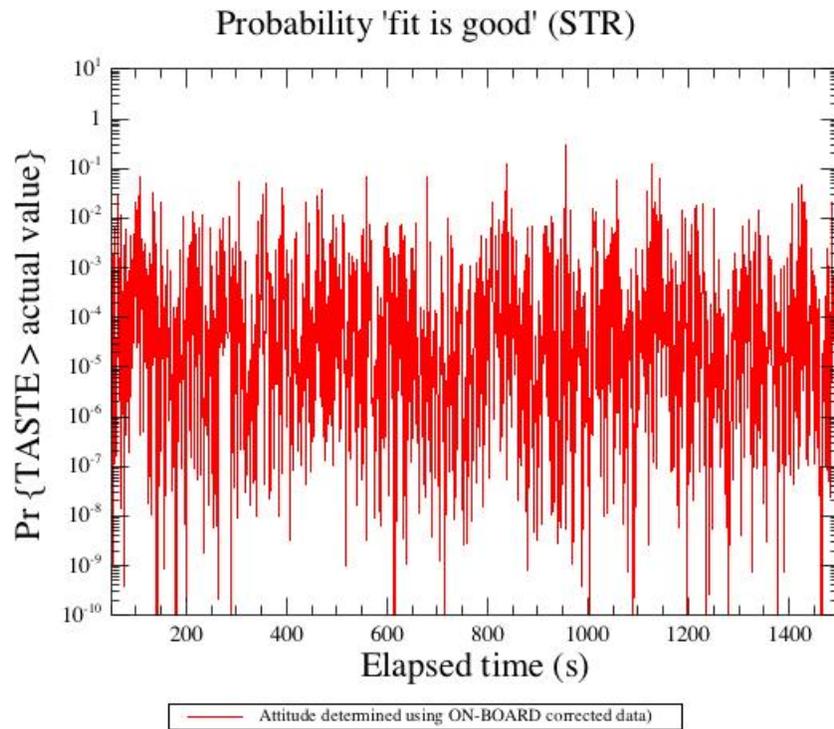


Figure 11: Quality of uncorrected attitude measurements – 1342197884



Column	Name	Type	Units	Description
1	obt	Double1d(*)	μs	On-board time.
2	newmeas_q	Double2d(*, 4)	None	New (corrected) measurement of ACA-frame attitude quaternion.
3	newmeas_ra	Double1d(*)	deg.	New (corrected) measurement of right ascension of ACA-frame x -axis.
4	newmeas_dec	Double1d(*)	deg.	New (corrected) measurement of declination of ACA-frame x -axis.
5	newmeas_roll	Double1d(*)	deg.	New (corrected) measurement of roll about ACA-frame x -axis.
6	newstr_q	Double2d(*, 4)	None	New (corrected) measurement of STR-frame attitude quaternion.
7	newstr_ra	Double1d(*)	deg.	New (corrected) measurement of right ascension of STR-frame x -axis.
8	newstr_dec	Double1d(*)	deg.	New (corrected) measurement of declination of STR-frame x -axis.
9	newstr_roll	Double1d(*)	deg.	New (corrected) measurement of roll about STR-frame x -axis.
10	TASTE	Double1d(*)	None	Value of TASTE variable from attitude determination.
11	prob_taste	Double1d(*)	None	Probability, p_{taste} , of such a large value of TASTE occurring at random; see p. 14, eq. (26).
12	num_bad	Int1d(*)	None	Number of stars excluded from attitude determination.
13	badStars	Int2d(*, 2)	None	IDs of excluded stars (if any).
14	nStarsUsed	Int1d(*)	None	Number of stars used in attitude determination.
15	sigma_x	Double1d(*)	arcsec.	Standard deviation of error about ACA-frame x -axis.
16	sigma_y	Double1d(*)	arcsec.	Standard deviation of error about ACA-frame y -axis.
17	sigma_z	Double1d(*)	arcsec.	Standard deviation of error about ACA-frame z -axis.
18	rho_yz	Double1d(*)	None	Correlation coefficient between errors about ACA-frame y - and z -axes.



19	rho_xz	Double1d(*)	None	Correlation coefficient between errors about ACA-frame x - and z -axes.
20	rho_xy	Double1d(*)	None	Correlation coefficient between errors about ACA-frame x - and y -axes.
21 (debug)	obmeas_q	Double2d(*, 4)	None	On-board measurement of ACA-frame attitude quaternion.
22 (debug)	obmeas_ra	Double1d(*)	deg.	On-board measurement of right ascension of ACA-frame x -axis.
23 (debug)	obmeas_dec	Double1d(*)	deg.	On-board measurement of declination of ACA-frame x -axis.
24 (debug)	obmeas_roll	Double1d(*)	deg.	On-board measurement of roll about ACA-frame x -axis.
25 (debug)	oldmeas_q	Double2d(*, 4)	None	Uncorrected measurement of ACA-frame attitude quaternion. ⁵⁵
26 (debug)	oldmeas_ra	Double1d(*)	deg.	Uncorrected measurement of right ascension of ACA-frame x -axis.
27 (debug)	oldmeas_dec	Double1d(*)	deg.	Uncorrected measurement of declination of ACA-frame x -axis.
28 (debug)	oldmeas_roll	Double1d(*)	deg.	Uncorrected measurement of roll about ACA-frame x -axis.
29 (debug)	rot_ob_q	Double2d(*, 4)	None	Quaternion giving rotation of new attitude measurement with respect to on-board attitude measurement.
30 (debug)	rot_ob_x	Double1d(*)	arcsec.	ACA-frame x -axis component of new attitude measurement with respect to on-board attitude measurement.
31 (debug)	rot_ob_y	Double1d(*)	arcsec.	ACA-frame y -axis component of new attitude measurement with respect to on-board attitude measurement.
32 (debug)	rot_ob_z	Double1d(*)	arcsec.	ACA-frame z -axis component of new attitude measurement with respect to on-board attitude measurement.
33 (debug)	rot_old_q	Double2d(*, 4)	None	Quaternion giving rotation of new attitude measurement with respect to uncorrected attitude measurement.

⁵⁵‘Uncorrected’ means computed on-ground using the measured star vectors in (6).



34 (debug)	rot_old_x	Double1d(*)	arcsec.	ACA-frame x -axis component of new attitude measurement with respect to uncorrected attitude measurement.
35 (debug)	rot_old_y	Double1d(*)	arcsec.	ACA-frame y -axis component of new attitude measurement with respect to uncorrected attitude measurement.
36 (debug)	rot_old_z	Double1d(*)	arcsec.	ACA-frame z -axis component of new attitude measurement with respect to uncorrected attitude measurement.
37 (debug)	TASTE_old	Double1d(*)	None	Value of TASTE variable from attitude determination giving oldmeas_q.
38 (debug)	prob_taste_old	Double1d(*)	None	Probability, p_{taste} , of such a large value of TASTE (that in TASTE_old) occurring at random; see p. 14, eq. (26).

Table 3: Contents of strAttitude



Column	Name	Type	Units	Description
1	obt	Double1d(*)	μs	On-board time, t_k , of gyro measurement.
2	gyratt_q	Double2d(*, 4)	None	Reconstructed ACA-frame attitude quaternion, $\hat{\mathbf{q}}_{\text{aca}}(t_k)$.
3	gyratt_ra	Double1d(*)	deg.	Reconstructed right ascension of ACA-frame x -axis.
4	gyratt_dec	Double1d(*)	deg.	Reconstructed declination of ACA-frame x -axis.
5	gyratt_roll	Double1d(*)	deg.	Reconstructed roll about ACA-frame x -axis.
6	prob_x	Double1d(*)	None	Quality associated with x -axis fit, i.e. $\Pr\{\min\{\chi_x^2\} > \chi_x^2(\hat{b}_x, \hat{c}_x)\}$, see eq. (47), p. 23.
7	prob_y	Double1d(*)	None	Quality associated with y -axis fit, i.e. $\Pr\{\min\{\chi_y^2\} > \chi_y^2(\hat{b}_y, \hat{c}_y)\}$, see eq. (47), p. 23.
8	prob_z	Double1d(*)	None	Quality associated with z -axis fit, i.e. $\Pr\{\min\{\chi_z^2\} > \chi_z^2(\hat{b}_z, \hat{c}_z)\}$, see eq. (47), p. 23.
9	sigma_x	Double1d(*)	arcsec.	Standard deviation of error in gyro-based reconstructed attitude about ACA-frame x -axis, i.e. $\sqrt{E[\tilde{\theta}_{x,k}^2]}$, see eq. (53), p. 24.
10	sigma_y	Double1d(*)	arcsec.	Standard deviation of error in gyro-based reconstructed attitude about ACA-frame y -axis, i.e. $\sqrt{E[\tilde{\theta}_{y,k}^2]}$, see eq. (53), p. 24.
11	sigma_z	Double1d(*)	arcsec.	Standard deviation of error in gyro-based reconstructed attitude about ACA-frame z -axis, i.e. $\sqrt{E[\tilde{\theta}_{z,k}^2]}$, see eq. (53), p. 24.
12 (debug)	maxrot_ref	Double1d(*)	deg.	Maximum rotation—over the interval used for the estimation—of the star tracker measurements



				from the reference attitude.
13 (debug)	num_meas	Int1d(*)	None	Number of measurements used in the fit (for each axis).
14 (debug)	drift_x	Double1d(*)	arcsec./s	Estimated (gyro) drift rate about ACA-frame <i>x</i> -axis.
15 (debug)	drift_y	Double1d(*)	arcsec./s	Estimated (gyro) drift rate about ACA-frame <i>y</i> -axis.
16 (debug)	drift_z	Double1d(*)	arcsec./s	Estimated (gyro) drift rate about ACA-frame <i>z</i> -axis.
17 (debug)	sig_dr_x	Double1d(*)	arcsec./s	Standard deviation of error in estimated (gyro) drift rate about ACA-frame <i>x</i> -axis.
18 (debug)	sig_dr_y	Double1d(*)	arcsec./s	Standard deviation of error in estimated (gyro) drift rate about ACA-frame <i>y</i> -axis.
19 (debug)	sig_dr_z	Double1d(*)	arcsec./s	Standard deviation of error in estimated (gyro) drift rate about ACA-frame <i>z</i> -axis.

Table 4: Contents of gyroAttitude



4 Summary of known issues

Below is a summary of the issues known to affect the versions of the main ‘pointing toolbox’ functions contained in build 3342 of HIPE 14.0: `CalcAttitudeTask` (version 1.10), `calcStrAttitude` (version 1.37) and `calcGyroAttitude` (version 1.31).⁵⁶ For further details, click on the ‘issue key’ to follow the link to the JIRA page describing the issue.

Table 5 provides an overview of the issues and attempts to classify them according to the likelihood of the User encountering the problem and the severity of the problem if encountered.

Issue key	Likelihood of problem	Severity
HCSS-...	Unlikely (0), Probable (1), Very likely (2), Always (3)	No impact (0), Minor (1), Moderate (2), Major (3)
19121	0–1	2
19122	1	2
19200	3	0
19201	0–1	3
20065	2	1–2

Table 5: Overview of issues

HCSS-19121

The Moore–Penrose pseudoinverse, G^+ , of the gyro alignment matrix (see p. 18) is being used in (39) to perform a least-squares fit and convert the four integrated gyro rates into three small-angle rotations. The function `calcGyroAttitude` currently does not check to see whether the output from the four gyros is consistent.

HCSS-19122

The software which is used to compute the star tracker CCD distortion maps compares the measured coordinates of each detected star with its expected coordinates. However, instead of using the best estimate of the spacecraft attitude to transform the catalogue coordinates of each star from the inertial reference frame to the Boresight Reference Frame (BRF), the software which produced the maps currently in use by the

⁵⁶This build was delivered on 22 December 2015, but the last change to the gyro-based pointing reconstruction software was made on 1 October 2015.



pointing reconstruction software used the attitude estimated using the ‘reference period’ parameters (i.e. the parameters which were on-board between ODs 320–762). The map generation software was modified in July 2014 to use an iterative approach and new distortion maps were generated [see 10]. However, these maps are still not being used by the pointing reconstruction software.

HCSS-19200

The creation date added by `calcAttitude` to the meta data of the new pointing product is not the date when the product was created.

HCSS-19201

The function `calcAttitude` overwrites the fields `filterQuat`, which contain the attitude quaternions from the on-board filter, with the attitude quaternions from the ground-based attitude reconstruction. Since the OBTs in the pointing product are left unmodified and since they do not correspond exactly with those found in the output table from `calcGyroAttitude`, each quaternion is overwritten with the quaternion corresponding to the closest matching time. Typically, the times match to within a few microseconds, but, as no test is performed, the mismatch may be larger. (Times differing by as much as 1.5 seconds have been observed.)

HCSS-20065

A suggestion to improve the performance of a particularly slow block of code in `calcStrAttitude`.



A Inversion of distortion correction equations

A.1 Using a modified Newton method

A simple and very effective way of inverting the system of algebraic equations (11), which avoids completely the necessity of having to estimate new sets of coefficients for the distortion correction polynomials, is to set $y_0 = y'$, $z_0 = z'$ and to iterate using:⁵⁷

$$\begin{aligned} y_{i+1} &= y' + y_i - F_1(y_i, z_i; k_0, \dots, k_7), \\ z_{i+1} &= z' + z_i - F_1(z_i, y_i; h_0, \dots, h_7). \end{aligned} \tag{54}$$

Writing (11) as

$$\begin{aligned} g_1(y, z) &\equiv F_1(y, z; k_0, \dots, k_7) - y' = 0, \\ g_2(y, z) &\equiv F_1(z, y; h_0, \dots, h_7) - z' = 0, \end{aligned} \tag{55}$$

Newton's method gives:

$$\begin{aligned} \begin{pmatrix} y_{i+1} \\ z_{i+1} \end{pmatrix} &= \begin{pmatrix} y_i \\ z_i \end{pmatrix} - \left[\frac{\partial(g_1, g_2)}{\partial(y, z)} \right]_{(y_i, z_i)}^{-1} \begin{pmatrix} g_1(y_i, z_i) \\ g_2(y_i, z_i) \end{pmatrix} \\ &= \begin{pmatrix} y_i \\ z_i \end{pmatrix} - \left[\begin{array}{cc} J_{11} & J_{12} \\ J_{21} & J_{22} \end{array} \right]_{(y_i, z_i)}^{-1} \begin{pmatrix} F_1(y_i, z_i; k_0, \dots, k_7) - y' \\ F_1(z_i, y_i; h_0, \dots, h_7) - z' \end{pmatrix}, \end{aligned} \tag{56}$$

where

$$\begin{aligned} J_{11} &= k_1 - 2k_5y - k_6z + k_3(3y^2 + z^2) + k_4(5y^4 + 6y^2z^2 + z^4), \\ J_{12} &= k_2 - k_6y - 2k_7z + 2k_3yz + 4k_4yz(y^2 + z^2), \\ J_{21} &= h_2 - 2h_7y - h_6z + 2h_3yz + 4h_4yz(y^2 + z^2), \\ J_{22} &= h_1 - h_6y - 2h_5z + h_3(y^2 + 3z^2) + h_4(y^4 + 6y^2z^2 + 5z^4). \end{aligned} \tag{57}$$

We see therefore that the intuitive scheme (54) corresponds to the approximation $J_{11} = J_{22} = 1$, $J_{12} = J_{21} = 0$. In general, the quadratic convergence exhibited by Newton's method is lost when one approximates the (inverse) Jacobian by a constant matrix and only linear convergence can be expected [12, pp. 109–110]. However, it appears that in the problem under consideration the approximation is sufficiently accurate that little of the quadratic

⁵⁷For the sake of clarity the subscript 'r' on y_r , z_r , y'_r and z'_r has been dropped here.



convergence is lost. That is, it has been found that two iterations of (54) are sufficient to match the accuracy of the current method and that a third iteration reduces the errors to below 0.0001".

A.2 Using power series

Berrighi claims that the solution of (1), or equivalently for our purposes (11), will exist in the form of Taylor series and that the "same level of accuracy" can be obtained by truncating these series at degree five [4, p. 84].⁵⁸

In an attempt to justify this, rewrite (11) as:

$$g_i(y, z; y', z') = 0, \quad i = 1, 2.$$

The implicit functions g_1 and g_2 are polynomials and hence analytic in y, z, y' and z' . Equations (57) show the value of the Jacobian determinant at the origin to be

$$\left| \frac{\partial(g_1, g_2)}{\partial(y, z)} \right|_{y=z=0} = |k_1 h_1 - k_2 h_2| \approx |k_1 h_1|.$$

Since this determinant is non-zero, as may be easily verified from the values found for the distortion correction coefficients, and the solution satisfies the conditions

$$g_i(0, 0; k_0, h_0) = 0, \quad i = 1, 2,$$

it follows from the theory of analytic implicit functions [e.g. 15, ch. 6] that y and z may be solved as power series of the form

$$\begin{aligned} y &= \sum_{i,j=0}^{\infty} c_{ij} (y' - k_0)^i (z' - h_0)^j, \\ z &= \sum_{i,j=0}^{\infty} d_{ij} (y' - k_0)^i (z' - h_0)^j, \end{aligned} \tag{58}$$

where the c_{ij} and d_{ij} are real constants and $c_{00} = d_{00} = 0$.

A formal solution may be obtained by the method of undetermined coefficients and Cauchy's method of majorants may then be used to show that

⁵⁸The second statement is understood to mean that the errors incurred in truncating the power series solution at the terms of degree five is negligible compared with the accuracy of the distortion correction model (a system of polynomials of degree five).



these series converge for sufficiently small $y' - k_0$ and $z' - h_0$. However, without further analysis, there is no guarantee that these series converge over the entire range of values corresponding to the star tracker's CCD and Berrighi's claim regarding the level of accuracy that may be achieved by truncating the series at the terms of degree five would seem to be entirely unjustified.



B Ignoring the inter-axis correlation of the measurement errors

The attitude measurement errors about the spacecraft x -axis are known to be highly correlated with those about the y - and z -axes. The effect of ignoring this correlation and using the decoupled equations (45) has been investigated using the data set for observation 1342197884 at on-board time 1654553791.038719 s. This is a data set for which the fits were found to be of good quality; see Figures 12–14.

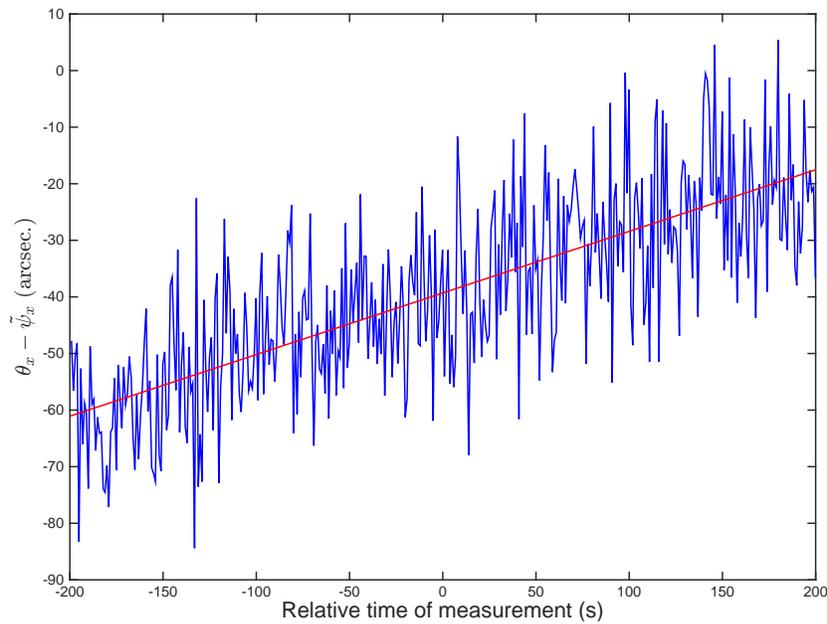


Figure 12: Least-squares fit of x -axis measurements

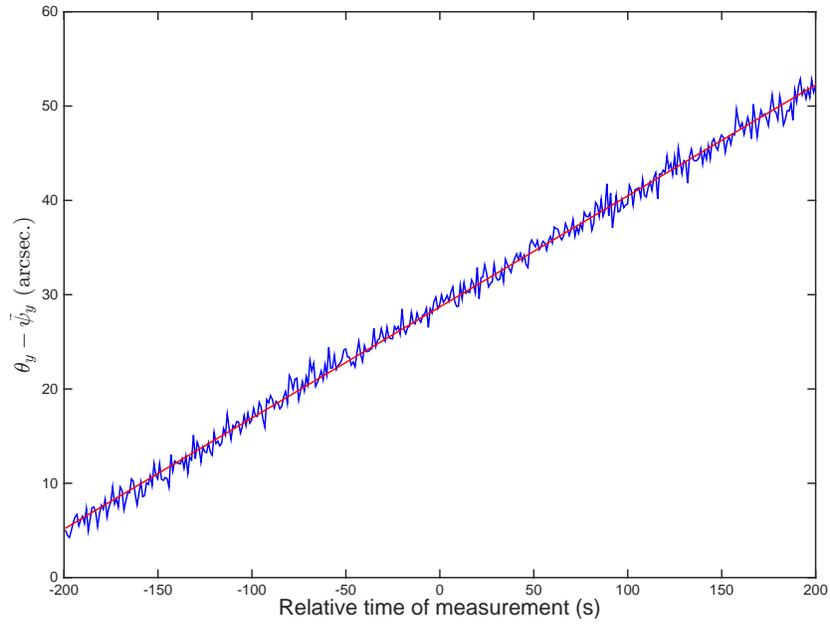


Figure 13: Least-squares fit of y -axis measurements

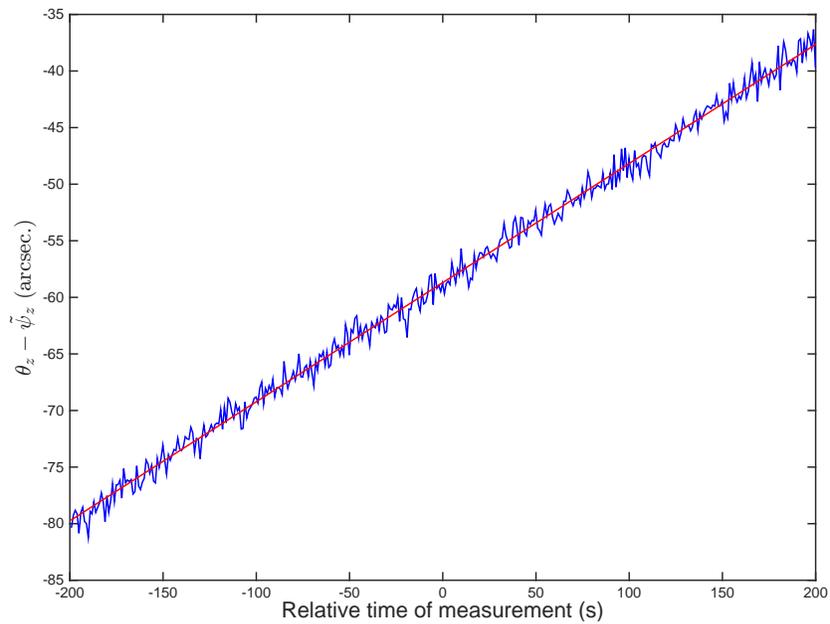


Figure 14: Least-squares fit of z -axis measurements



(a) Attitude estimate

Table 6 compares the estimates of the three parameters c_x , c_y and c_z .⁵⁹ It can

	Full equations	Decoupled equations
\hat{c}_x	-39.304118	-39.304121
\hat{c}_y	28.700927	28.700929
\hat{c}_z	-58.697253	-58.697259

Table 6: Estimates of c_x , c_y and c_z (arcseconds).

be seen that the use of the decoupled equations has a negligible effect on the estimated attitude. For each of the three spacecraft axes the simplification introduces an error of less than 0.00001".

(b) Confidence regions

As noted in Section 2.2.4, for the decoupled problem, the expected variances of the errors in c_x , c_y and c_z are provided by $(X_x^T X_x)^{-1}[2, 2]$, \dots , $(X_z^T X_z)^{-1}[2, 2]$. For the full, coupled problem we may obtain the full covariance matrix for these errors by extracting the even rows and columns of:

$$E[(\beta - \hat{\beta})(\beta - \hat{\beta})^T] = [(X^*)^T X^*]^{-1} = [X^T \Omega^{-1} X]^{-1}.$$

For the dataset in question, the values so obtained are shown in Table 7.

	Full equations	Decoupled equations
$E[\tilde{c}_x^2]$	0.331572324300	0.331572325830
$E[\tilde{c}_y^2]$	0.002357191189	0.002357191195
$E[\tilde{c}_z^2]$	0.002749527116	0.002749527124
$E[\tilde{c}_y \tilde{c}_z]$	-0.000054237934	–
$E[\tilde{c}_x \tilde{c}_z]$	0.011613113893	–
$E[\tilde{c}_x \tilde{c}_y]$	-0.001651280255	–

Table 7: Elements of formal covariance matrix (arcseconds²).

To obtain confidence regions for the attitude errors (about the y and z spacecraft axes) we project the regions in β -space defined by $J(\beta) = \text{constant}$

⁵⁹The estimates of b_x , b_y and b_z are of little interest as they are not used to reconstruct the attitude.



on to the $c_y c_z$ -plane.⁶⁰ If the measurement errors are normally distributed (see Appendix C), then the resulting elliptical boundaries for the attitude errors $\tilde{\theta}_y$ and $\tilde{\theta}_z$ are related to the formal covariance matrix by:⁶¹

$$(\tilde{\theta}_y \quad \tilde{\theta}_z) \begin{pmatrix} E[\tilde{c}_y^2] & E[\tilde{c}_y \tilde{c}_z] \\ E[\tilde{c}_y \tilde{c}_z] & E[\tilde{c}_z^2] \end{pmatrix}^{-1} \begin{pmatrix} \tilde{\theta}_y \\ \tilde{\theta}_z \end{pmatrix} = \Delta, \quad (59)$$

where the value of Δ depends on the confidence level (e.g. $\Delta = 2.30$ for a 68.3% level and $\Delta = 11.8$ for a 99.7% level). Writing

$$\begin{pmatrix} E[\tilde{c}_y^2] & E[\tilde{c}_y \tilde{c}_z] \\ E[\tilde{c}_y \tilde{c}_z] & E[\tilde{c}_z^2] \end{pmatrix}^{-1} = P^{-1} D P,$$

where $P = \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix}$ and $D = \begin{pmatrix} d_1 & 0 \\ 0 & d_2 \end{pmatrix}$, we find (using the values in Table 7) that $\alpha = 7.73^\circ$, $d_1 = 425.6 \text{ arcsec.}^{-2}$ and $d_2 = 362.7 \text{ arcsec.}^{-2}$ for the full equations and that $\alpha = 0$, $d_1 = 424.2 \text{ arcsec.}^{-2}$ and $d_2 = 363.7 \text{ arcsec.}^{-2}$ for the decoupled equations. That is, the error ellipse is approximately circular (it has an eccentricity of just 0.04) and the result of using the decoupled equations, as opposed to the full equations, introduces errors in the directions and magnitudes of its principal axes of approximately 7.7° and 0.15% respectively.⁶² (At the 99.7% confidence level, the error ellipse has principal axes of lengths $0.33''$ and $0.36''$.)

(c) Goodness-of-fit

For the full problem, there are $3(n_s - n_d)$ equations and 6 estimated parameters. Therefore the p -value associated with the goodness-of-fit is given

⁶⁰Since the instrument boresights are aligned closely with the spacecraft x -axis, the attitude error about the x -axis will have little effect on the pointing error. To include this error we would instead project on to the $c_x c_y c_z$ subspace.

⁶¹See [16, pp. 690–693].

⁶²Setting $\begin{pmatrix} \tilde{\theta}'_y \\ \tilde{\theta}'_z \end{pmatrix} = P \begin{pmatrix} \tilde{\theta}_y \\ \tilde{\theta}_z \end{pmatrix}$, (59) may be rewritten as $\frac{d_1}{\Delta} \tilde{\theta}'_y{}^2 + \frac{d_2}{\Delta} \tilde{\theta}'_z{}^2 = 1$, showing that the ellipse has principal axes of lengths $2\sqrt{\frac{\Delta}{d_1}}$ and $2\sqrt{\frac{\Delta}{d_2}}$. Assuming without loss of generality that $d_1 > d_2$, the ellipse has eccentricity $\frac{\sqrt{d_1/d_2} - 1}{\sqrt{d_1/d_2} + 1}$.



by:

$$p_{\text{full}} \equiv \Pr\{\min\{J\} > J(\hat{\beta})\} = 1 - P\left(\frac{3(n_s - n_d) - 6}{2}, \frac{J(\hat{\beta})}{2}\right), \quad (60)$$

where P is the lower (regularized) incomplete gamma function. For the decoupled equations, the p -values, p_x , p_y and p_z , associated with the three least-squares fits are obtained from (47) and then combined according to (48) and (49). For our test dataset, the computed p -values are shown in Table 8. Taking into account that the least-squares problems being solved are slightly different, the agreement between the combined p -value and that for the full problem is acceptable.

	x -axis	y -axis	z -axis	Combined (Fisher)	Full
p -value	0.557	0.965	1.000	0.975	1.000

Table 8: Goodness-of-fit p -values.

C Normality of the measurement errors

To investigate the assumption that the measurement errors are normally distributed we look instead at the distribution of the normalized residuals for the dataset used in Appendix B.⁶³ Figures 15–17 compare the empirical (cumulative) distribution functions, S_x , S_y and S_z , of the normalized residuals from each of the three single-axis fits, with the distribution function, F , for the standard normal distribution.

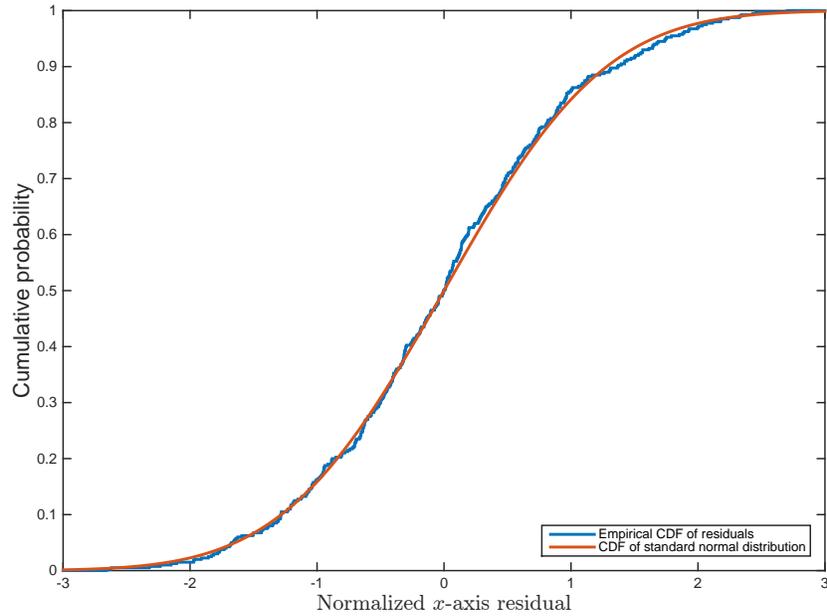


Figure 15: CDF of normalized x -axis residuals

To test the assumption (our null hypothesis) that the normalized residuals were drawn from a standard normal distribution, we calculate the Kolmogorov-Smirnov statistic

$$D = \max_{-\infty < u < \infty} |S_r(u) - F(u)|, \quad r \in \{x, y, z\},$$

where $F(u) = [1 + \text{erf}(u/\sqrt{2})]/2$. Assuming the null hypothesis to be true, the probability of obtaining a value of the test statistic which exceeds the

⁶³Each residual is normalized by dividing it by the standard deviation of the expected error in the corresponding measurement.

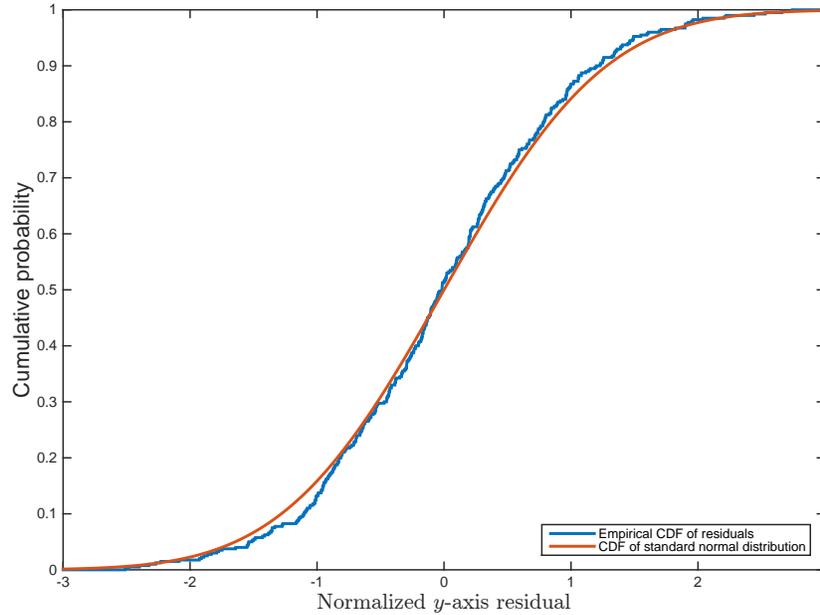


Figure 16: CDF of normalized y -axis residuals

observed value, D_{obs} , may be approximated by:

$$\Pr\{D > D_{\text{obs}}\} = 2 \sum_{j=1}^{\infty} (-1)^{j-1} e^{-2j^2(D^*)^2},$$

where $D^* = (\sqrt{N} + 0.12 + 0.11/\sqrt{N}) D$ and $N = 400$ is the number of residuals.⁶⁴ The results, shown in Table 9, lead us to accept, at the 5% significance level, the hypothesis that the residuals are normally distributed. However, as it is known that the Kolmogorov-Smirnov test is fairly insensitive

Axis	K-S statistic D_{obs}	$\Pr\{D > D_{\text{obs}}\}$
x	0.035	0.71
y	0.039	0.57
z	0.064	0.07

Table 9: Results from Kolmogorov-Smirnov test.

to deviations from $F(u)$ which occur away from the median value of u , it

⁶⁴See [16, pp. 617–620]. The modified test statistic D^* was first introduced in [22].

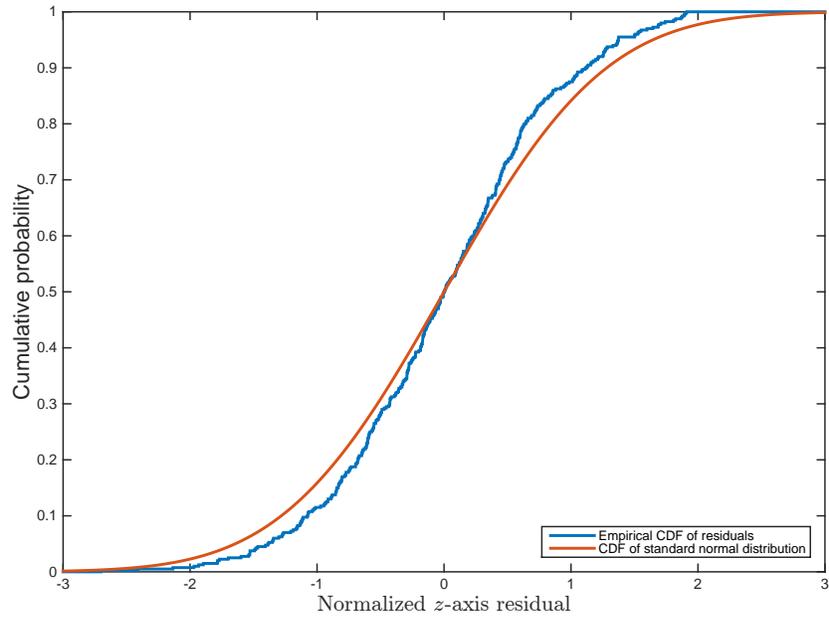


Figure 17: CDF of normalized z -axis residuals

might be advisable to repeat the test using an alternative statistic, such as Kuiper's.



References

- [1] ACMS Team (Dutch Space). Herschel-Planck ACMS User Manual. Technical Report H-P-4-DS-MA-001, Issue 4/4, Dutch Space, Feb 2009.
- [2] H. Aussel. Towards an accurate reconstruction of Herschel pointing. Technical Report SAp-PACS-HA-0729-11, Issue 1.0, CEA Saclay, Nov 2011.
- [3] H. Aussel. PACS pointing jitter : the GYR view. Technical Report SAp-PACS-HA-0730-12, Issue 1.0, CEA Saclay, Aug 2012.
- [4] G. Berrighi. ASTR for HERSCHEL/PLANCK User Manual. Technical Report H-P-4-GAF-MA-0001, Issue 7, Galileo Avionica, Oct 2006.
- [5] G. Berrighi. ASTR for HERSCHEL/PLANCK Communications Interface Control Document. Technical Report H-P-4-GAF-IC-0001, Issue 7, Galileo Avionica, 2006.
- [6] J. Bradley. A Letter from the Reverend Mr. James Bradley Savilian Professor of Astronomy at Oxford, and F.R.S. to Dr. Edmond Halley Astronom. Reg. &c. giving an Account of a new discovered Motion of the Fix'd Stars. *Philosophical Transactions of the Royal Society*, 35: 637–661, Jan 1727.
- [7] H. Feuchtgruber. Herchel Guide Stars. Technical Report PICC-ME-TN-040, Issue 1.0, Herschel PACS, Dec 2011.
- [8] H. Feuchtgruber. Herschel STR-A CCD Sub-Pixel Structure. Technical Report PICC-ME-TN-041, Issue 2.0, Herschel PACS, Mar 2012.
- [9] H. Feuchtgruber. Reconstruction of the Herschel Pointing Jitter. Technical Report PICC-ME-TN-042, Draft, Revision 1.0, Herschel PACS, Sep 2012.
- [10] H. Feuchtgruber. Herschel STR-A CCD Sub-Pixel Structure. Technical Report PICC-ME-TN-041, Issue 4.0, Herschel PACS, Jul 2014.
- [11] R. A. Fisher. *Statistical Methods for Research Workers*. Oliver and Boyd, fifth edition, 1934.



- [12] R. Kress. *Numerical Analysis*. Number 181 in Graduate Texts in Mathematics. Springer, 1998.
- [13] D. Lutz. Some tests for `toff_gyro` in `calcAttitude`. Oct 2014.
- [14] F. L. Markley and D. Mortari. Quaternion Attitude Estimation Using Vector Observations. *Journal of the Astronautical Sciences*, 48(2–3): 359–380, Apr–Sep 2000.
- [15] F. R. Moulton. *Differential Equations*. The Macmillan Company, 1930. Reprinted by Dover Publications, New York, 1958.
- [16] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in Fortran 77: The Art of Scientific Computing*. Cambridge University Press, second edition, 1992. Reprinted with corrections, 1997.
- [17] M. Sánchez-Portal. Herschel Pointing Product Specification. Technical Report HERSCHEL-HSC-DOC-0662, Issue 1.14, Herschel Science Centre, Nov 2014.
- [18] M. D. Shuster. Maximum Likelihood Estimation of Spacecraft Attitude. *Journal of the Astronautical Sciences*, 37(1):79–88, Jan–Mar 1989.
- [19] M. D. Shuster. The Generalized Wahba Problem. *Journal of the Astronautical Sciences*, 54(2):245–259, Apr–Jun 2006.
- [20] M. D. Shuster and D. C. Freesland. The Statistics of TASTE and the Inflight Estimation of Sensor Precision. The NASA/GSFC Flight Mechanics Symposium, Oct 2005.
- [21] M. D. Shuster and S. D. Oh. Three-Axis Attitude Determination from Vector Observations. *Journal of Guidance and Control*, 4(1):70–77, Jan–Feb 1981.
- [22] M. A. Stephens. Use of the Kolmogorov-Smirnov, Cramer-Von Mises and Related Statistics Without Extensive Tables. *Journal of the Royal Statistical Society, Series B (Methodological)*, 32(1):115–122, 1970.
- [23] M. Tuttlebee. HERSCHEL/PLANCK Star Tracker Focal Length Assessment Using Raw Star Telemetry Data. Technical Report PT-CMOC-FD-TN-2724-OPS-GFE, Issue 1.0, ESOC Flight Dynamics, Mar 2011.



- [24] M. Tuttlebee. HERSCHEL/PLANCK Star Tracker Performance Assessment and Calibration. Technical Report PT-CMOC-OPS-RP-6435-HSO-GF, Issue 1.0, ESOC Flight Dynamics, Aug 2013.
- [25] G. Wahba. A Least Squares Estimate of Satellite Attitude. *SIAM Review*, 7(3):409, Jul 1965.
- [26] J. R. Wertz, editor. *Spacecraft Attitude Determination and Control*. Kluwer Academic Publishers, 1978. Reprinted in 1988.