15. REAL-TIME ATTITUDE DETERMINATION WITH TWO GYROS

Radiation damage was considered to be responsible for the progressive failure of the gyroscopes. In the nominal mode of operation these gyros were essential for acquisition and maintenance of attitude knowledge and control. In principle the satellite itself, with adequately known external and internally-actuated torques, allows the integration of its attitude evolution with less noise than integration of the gyro outputs. However, this would require both a very accurate starting point and very sophisticated modelling. The accuracy of such a description of attitude evolution could only be maintained through the observation of star transits across the star mapper grids. The time between these events must be sufficiently short for the integration to maintain the required accuracy. In the originally intended geostationary orbit it seems likely that this notion would have provided the incentive to operate with the gyros turned off, thus suppressing jitter and parasitic torque. In the actual geostationary transfer orbit the long interruptions of star transit observations coincided with the least known, yet highest perturbing torques, which would inevitably cause loss of attitude to the point of requiring reinitialisation. Thus for a less than three-gyro mode of operation not only a model to replace the real-time onboard input of the absent gyro was needed, but also a much faster, and thus more sophisticated attitude determination and initialisation on ground. As redundancy in the gyroscopes disappeared, a substantial effort was put into developing the procedures to work with only two gyroscopes, and eventually without gyroscopes altogether. The two-gyro mode was fully implemented, and eventually used for the last months of scientific observations.

15.1. Two-Gyro Operations Development History

In June 1990, gyro 4, which monitored the scan rate about the principal spin axis, failed. It was replaced in the attitude control loop by the backup unit, gyro 5. Before the failure, gyro 4 had exhibited a large and fluctuating noise envelope (exceeding 1 arcsec s^{-1} amplitude at times). Subsequent monitoring showed that gyro 5 exhibited similar noise characteristics to those of the failed gyro.

Given the repeated apparent run-downs of gyro 4 and the anomalous noise properties of gyro 5, there was a high risk that both could fail irreparably. In order to continue

the mission without any z-gyro (the z referring to the z-axis in the spacecraft reference system used for real-time attitude determination, as described in Chapter 14), a series of studies were performed. They resulted in an attitude control software modification, called the 'two-gyro patch', and a number of ground software changes/updates which allowed the continuation of the mission using only two gyros.

The revised on-board real-time attitude determination algorithm replaced the z-gyro input by a simple sinusoidal model for the disturbance torque around the z-axis near apogee, combined with a revised Kalman filter to estimate not only the z-axis Tait-Bryan error angle (as for the three-gyro real-time attitude determination) but also the difference between the nominal and actual spin rate and the acceleration around the z-axis. Both the Kalman filter state vector and the sinusoidal model had to be initialised from ground after every perigee using the revised ground real-time attitude determination software. The on-board attitude control logic, indicating the changes necessary in going from three-gyro to two-gyro operations, are shown in Figure 14.2.

The patch was run successfully in parallel with the actual real-time attitude determination in order to validate both the on-board and ground modifications. These changes were extensively tested during the second half of 1991 and proven to function correctly.

On 10 July 1992, after continuous poor performance, gyro 5 was removed from the control loop and the two-gyro real-time attitude determination system was used to control the spacecraft for the first time. By 11 July 1992, after parameter tuning on-board and refinement of ground procedures, the first scientific data were obtained. The concept of two-gyro real-time attitude determination had been proven, and when converged, gave better performance than the three-gyro real-time attitude determination with the noisy gyro 5.

Although it had been proved that the ground software developed during 1991 and 1992 was adequate for initialising real-time attitude determination, it was far from optimal. Moreover, although giving generally better performance in the apogee region, the two-gyro real-time attitude determination suffered two distinct disadvantages:

(a) divergence after occultations was frequent, leading to lost science data before the ground real-time attitude determination software could re-initialise on-board real-time attitude determination;

(b) there was little control of the spin rate through perigee, leading to a lengthy procedure of real-time attitude determination initialisation after every perigee. This procedure took between 1.5 and 2 hours at that time, in comparison with 20–40 min with the three-gyro real-time attitude determination (with a high probability that real-time attitude determination would remain converged after perigee). Given that typically 9.5 hours was covered by ground station per orbit, a loss of 1 to 2 hours science data per orbit was considerable.

After an emergency sun reacquisition on 4 August 1992, in which the satellite spin axis was autonomously slewed to the Sun direction, a recovery in the performance of gyro 5 prompted a return to three-gyro real-time attitude determination operations on 8 August 1992. However, on 11 August after a scanning law acquisition manoeuvre to return to the nominal scanning law, gyro 5 failed for the last time causing an emergency sun reacquisition and a spin-up to 0.45 rpm. The decision was taken to extend the spacecraft spin-up phase as much as possible, i.e. until the increasing solar aspect angle reached the limit of 43°. There were two main reasons for doing so:

(a) the early operational experience with the two-gyro real-time attitude determination before the emergency sun reacquisition occurred, proved to be valuable in highlighting various improvements which could be made to speed up real-time attitude determination initialisation. While the spacecraft was spin-stabilised, it was safe enough to remove the flight dynamics team from operational shift work to daytime development of enhanced ground real-time attitude determination software, thereby increasing the speed of production of the software by a factor of 3 to 4. As stated above, these changes increased the speed and reliability of real-time attitude determination initialisation once the operations were resumed;

(b) it was hoped at that time that the mission lifetime would be dictated by the gas consumables. Given that no gas was being used in the spin-up phase, by delaying spin-down, one was actually extending the lifetime of the mission with a subsequent improvement in the final accuracies of the proper motions.

This approach was agreed between ESOC, ESTEC and the data reduction consortia. The spacecraft remained spun up until 1 October 1992, when the increasing sun aspect angle necessitated an additional manoeuvre to sun-pointing mode. By this time, most of the planned software changes had been completed. Full operations resumed with the new system, collecting science data for approximately four more months.

15.2. Operational Requirements

In running with two-gyro control, several new operational procedures had to be performed:

(a) the higher disturbance torques close to perigee caused a net angular acceleration around the rotation axis. These acceleration terms were difficult to model and predict as there was no ground station coverage possible at so low a height (500 – 600 km). Nominally, the scan rate was constrained by the on-board attitude control software to be within approximately 2 per cent of the nominal rate of 168.75 arcsec s^{-1}. Without the z-axis gyro input, much larger excursions were expected. As far as possible, these excursions had to be minimised through perigee;

(b) on receiving signal from the spacecraft after perigee it was necessary to measure the scan rate on-ground and correct it, if necessary, by direct thruster command;

(c) the spacecraft lost real-time attitude determination control through perigee. Therefore it was necessary to estimate on-ground the attitude, body rates, (particularly around the z-axis) and correct the on-board values in real-time. It was also necessary to be able to reset the drift values on the remaining gyros;

(d) most importantly, after re-initialisation, the revised real-time attitude determination software on-board had to be capable of remaining converged on the correct attitude estimate to allow scientific data to be collected, accurate attitude information being necessary for the image dissector tubes to pilot successfully.

On-Board Memory Constraints

Memory space on-board was limited. Some of the spare memory already provided additional gyro/thruster monitoring during loss of signal, preventing an irretrievable spin-up of the spacecraft should the second spin-rate gyro fail. A highly sophisticated model would not fit the available space that was left. The model that was adopted is described in detail in the next section.

The revised on-board software was devised to fit into the available spare memory and could be run in parallel with the nominal three-gyro system. This allowed, during the first half of 1992, the new software to be tested on-board in a monitoring mode, i.e. the attitude was estimated without using the (still) operational spin-axis gyro. However this information was merely dumped in the telemetry and was not used to control the spacecraft. Rather the nominal on-board system continued to control the satellite until the time that the z-axis gyro failed.

Ground System Requirements

In addition to the dramatic changes required to the on-board real-time attitude determination, the ground system also required significant modification.

Raw star mapper data were only processed by the telemetry pre-processing and star mapper filtering task in the ground real-time attitude determination chain (see Chapter 14). Therefore it had to perform the scan-rate estimation if no z-gyro was operational. This led to the following constraints on the scan-rate estimation algorithm:

(a) the algorithm performance had to comply with the real-time constraints of the ground attitude determination within the resource limitations imposed by the available hardware and other essential control tasks running in parallel;

(b) implementation of the scan-rate estimation algorithm had to involve minimum changes to the existing telemetry pre-processing and star mapper filtering software. Algorithms and techniques employed in the nominal algorithm had to be retained wherever possible and interfaces to the original software kept simple. This reduced problems of integrating and testing new software and minimised development time. This was important at the start of the development since the lifetime of the remaining z-gyro was unknown;

(c) the scan-rate estimation algorithm could not rely on *a priori* knowledge of either the expected transit time or detailed characteristics of the slit system being crossed, since this information was only derivable from subsequent tasks in the ground system;

(d) the interfaces to other tasks in the ground attitude system had to be unaltered wherever possible to minimise software changes.

The operational constraints imposed on the on-ground real-time attitude determination gave it distinctive problems in star recognition which were quite different from either the on-board real-time attitude determination system or the post-processing work performed by the scientific data reduction consortia. For example, on-board real-time attitude determination could not successfully recognise stars in the star mapper data unless it could calculate the spacecraft attitude accurately. Ground real-time attitude

determination had to identify stars, with little or no *a priori* information concerning the spacecraft rotational phase around the spin axis. Besides, ground real-time attitude determination had also to be able to operate under these new conditions with very much larger deviations in the spin rate. Ground real-time attitude determination had also to straddle the twin hurdles of providing sufficiently robust and accurate attitude determination, when the finite resources of the on-board system were unable to cope, and performing it in near real-time. This latter constraint was not imposed on the data reduction consortia who were able to re-constitute the attitude off-line to much higher accuracy, given the good *a priori* attitude information of the on-board system, that enabled the image dissector tube piloting for the acquisition of good scientific data.

15.3. On-Board Software

On-Board Real-Time Attitude Determination using Two Gyros

As discussed above, given that on-board memory was limited, changes to the real-time attitude determination algorithm were limited in scope. A full description of the design concepts of real-time attitude determination was given in Chapter 14. This section concentrates on the differences between the three- and two-gyro systems.

In the nominal on-board system, the state vector in the Kalman filter consisted of the three Tait-Bryan error angles and three gyro drifts. All six were defined with respect to a co-rotating cartesian coordinate system. The integrated gyro outputs were used to estimate the attitude evolution and the star mapper transit times were used as measurements to update the state vector.

Without a z-gyro, it was necessary to include the angular rate around the z-axis within the state vector and to allow that to be updated by the star mapper measurements. Studies by Matra Marconi Space showed that it was feasible to replace the z-gyro rate inputs by a dynamic model predicting ψ, the Tait-Bryan error angle around the z-axis. The vector would be increased to dimension 8 to include separately first and second-order terms in ψ and the z-gyro drift would be replaced by $d\psi/dt$ and $d^2\psi/dt^2$.

Consequently the revised state vector was:
- ϕ: Tait-Bryan error angle around the spacecraft x-axis;
- θ: Tait-Bryan error angle around the spacecraft y-axis;
- ψ_1: Tait-Bryan error angle around the spacecraft z-axis (first-order term);
- ψ_2: Tait-Bryan error angle around the spacecraft z-axis (second-order term);
- d_x: gyro drift around the spacecraft x-axis;
- d_y: gyro drift around the spacecraft y-axis;
- d_z: scan rate drift from nominal (168.75 arcsec s^{-1});
- γ_z: rate of scan rate drift.

Above 20 000 km where the radiation pressure was the dominant disturbance torque, the acceleration around the z-axis could be successfully modelled by a harmonic of three times the nominal scan-rate. Such a model already existed on-board in the 'normal mode' control software for use in optimising thruster on-times (see Chapter 13), but up to now had not been used for attitude determination. The time dependency of the model was provided by the variables $(S_{v3}(k), C_{v3}(k))$, which were already derived for

three-gyro operations (see Equation 13.2). The two new coefficients required to define the amplitude and phase of this oscillatory model (A_d, B_d) were stable over at least one orbit. With frequent on-ground calibration (a procedure carried out since the start of the mission) such a model was used by providing a separate correction term for ψ (denoted ψ_d) outside the Kalman filter, where (at time t_k):

$$\psi_d = A_d\ C_{v3}(k) + B_d\ S_{v3}(k) \qquad [15.1]$$

Therefore:

$$\psi = \psi_1 + \psi_2 + \psi_d \qquad [15.2]$$

Gyro data from two transverse gyros gave angular velocities (body rates) about the x and y axes. Gyro processing was performed using an updated gyro projection matrix to obtain estimates of the spacecraft body rates (G_x, G_y, G_z) around all three axes:

$$\begin{aligned} G_z &= \omega_Z\ t_i \\ G_x &= a_1 G_1 + b_1 G_2 + c_1 G_3 + d_1 G_z \\ G_y &= a_2 G_1 + b_2 G_2 + c_2 G_3 + d_2 G_z \end{aligned} \qquad [15.3]$$

where ω_Z is the nominal scanning law Z rate, (the z error-rate effect was negligible) and G_i ($i = 1, 2, 3$) are gyro increments over $t_i = 16/15$ s for the (x, y) plane gyros. The coefficients of non-selected gyros were set to zero by telecommand.

The spacecraft attitude and remaining gyro drift rates were predicted as follows:

$$\begin{pmatrix} \phi(k+1) \\ \theta(k+1) \\ d_x(k+1) \\ d_y(k+1) \end{pmatrix} = \begin{pmatrix} 1 & \omega_Z\ t_i & -t_i & 0 \\ -\omega_Z\ t_i & 1 & 0 & -t_i \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \phi(k) \\ \theta(k) \\ d_x(k) \\ d_y(k) \end{pmatrix} + \begin{pmatrix} G_x(k+1) - \omega_X(k)\ t_i - G_y(k+1)\ \psi(k) \\ G_y(k+1) - \omega_Y(k)\ t_i - G_x(k+1)\ \psi(k) \\ 0 \\ 0 \end{pmatrix} \qquad [15.4]$$

$$\begin{pmatrix} \psi_1(k+1) \\ \psi_2(k+1) \\ d_z(k+1) \\ \gamma_z(k+1) \end{pmatrix} = \begin{pmatrix} 1 & 0 & t_i & 0 \\ 0 & 1 & 0 & t_i^2/2 \\ 0 & 0 & 1 & t_i \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \psi_1(k) \\ \psi_2(k) \\ d_z(k) \\ \gamma_z(k) \end{pmatrix} + \begin{pmatrix} -G_x(k+1)\theta(k) \\ 0 \\ 0 \\ 0 \end{pmatrix} \qquad [15.5]$$

Spacecraft body rates around all three axes could therefore be estimated as:

$$\text{rate}(k+1) = (\text{attitude}(k+1) - \text{attitude}(k))/t_i \qquad [15.6]$$

In the event of actuation, for the following cycle, corrective terms were added to Equation 15.4 such that:

$$\begin{aligned} \psi_1(k+1) &= \psi_1(k) + d_z(k)\ t_i - g_x(k+1)\ \theta(k) + T_{az}\ t_{\text{on,z}}(t_i - t_{\text{on,z}}/2) \\ d_z(k+1) &= d_z(k) + \gamma_z(k)\ t_i + T_{az}\ t_{\text{on,z}} \end{aligned} \qquad [15.7]$$

where T_{az} is the mean thruster acceleration on-board parameter (fixed) over actuation, $t_{\text{on,z}}$ is the applied thruster on-time. Consequently, an accurate thruster calibration was required and this was performed and provided to ESOC by the data reduction consortia (see Chapter 9).

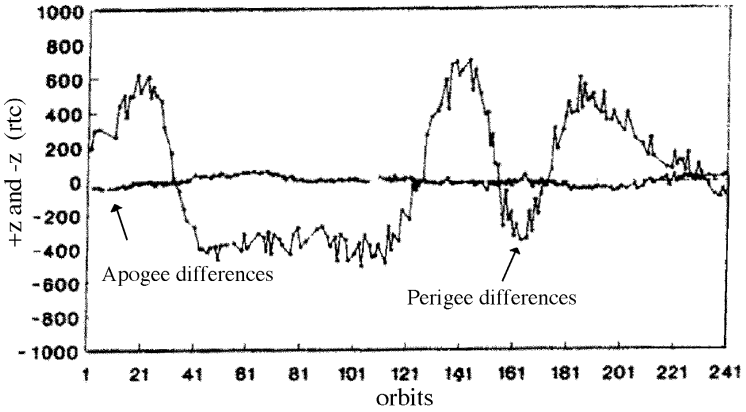

Figure 15.1. *z-thruster on-time differences (expressed in rtc units of 1/75 s) as accumulated around perigee and apogee for 240 orbits. See text for details.*

Measurement extrapolation and innovation (R_I) computation were performed in a similar manner to that used for three-gyro real-time attitude determination except that new fixed gain sets $\mathrm{KG}_{\mathrm{new}}$ were used. Attitude updates were consequently computed as:

$$X^{+}_{\mathrm{new}} = X_{\mathrm{new}} + R_I \ \mathrm{KG}_{\mathrm{new}}(\mathrm{slit}, \mathrm{quality}) \qquad [15.8]$$

where:

$$\begin{aligned} X_{\mathrm{new}} &= (\phi, \theta, \psi, d_x, d_y, x_z, \gamma_z) \\ \psi_1^{+} &= \psi^{+} - \psi_d \\ \psi_2^{+} &= 0 \end{aligned} \qquad [15.9]$$

Modification of the attitude for on-board computer extrapolation on cycles around actuation was performed as for three-gyro real-time attitude determination.

Perigee Scan-Rate Control

Through perigee passage, the complex disturbance torques around the z-axis were not modelled on-board to sufficient accuracy for attitude determination. In addition the accelerations could not be determined in the Kalman filter from star transit measurements because the star mapper electronics were blinded by noise from the van Allen belts.

It was soon realised that a further z-gyro failure could result in an irretrievable spin-up of the spacecraft. A software patch was implemented which allowed the accumulation of the $+z$ and $-z$ thruster on-times. Should either exceed a certain threshold which was reset every orbit, further firings of the z-thruster would be inhibited. This facility had the additional benefit of giving, on readout, information about the torques around the z-axis. The difference between the $+z$ and $-z$ total on-time allowed the determination of the total spin-up/down of the spacecraft if actuations about the z-axis were inhibited. Figure 15.1 shows a graph of on-time differences around perigee and apogee over 240 orbits (approximately 107 days). Thruster units (rtc) are in 1/75 s, which equated to a change in angular velocity of approximately 0.12 arcsec s^{-1} rtc^{-1}. Apogee differences always remained small. However perigee differences could be as great as 800 rtc units (100 arcsec s^{-1} change from the nominal rate of 168.75 arcsec s^{-1}).

No direct information was available about the likely phase angle drift around perigee, but drifts of several degrees were expected. To get more information on this aspect and provide a better overall understanding of the torques in this region, a further on-board software patch was later implemented, which allowed the storage and later dumping of telemetry parameters during loss of coverage. This 'spy patch' allowed ground to replay selected gyro and thruster information around the perigee region, giving more insight and more accurate predictions of scan-rate and scan angle deviations.

15.4. On-Ground Software

Scan-Rate Estimation and Correction

Ground real-time attitude determination, required a more sophisticated approach than that employed on-board, where the star mapper would not only measure the virtual transit time of a star as it crossed one of the sets of four slits, but would measure the spread in the four peak data stream to obtain the spin rate directly. This involved determining the best correlated signal as a function of the position and spread of the four-peak signal. Support was also required from ground in re-calibration and maintenance of the on-board harmonic model.

After the perigee passage the spacecraft would have an unknown scan rate about the z-axis. Once the ground station had acquired a signal from the spacecraft and the background noise had decreased to an acceptable level, the scan rate could be estimated. This was a time-critical task which directly affected the quantity of scientific data which could be collected in each orbit. It was an additional task to the standard operations and thus had to be performed as quickly as possible, but without sacrificing the accuracy of the achieved scan-rate after firing the thrusters.

The basis of the method was to estimate as accurately as possible the time taken to cross between the first and fourth slits. Since the distance between these slits had been accurately calibrated in orbit, this transit time allowed the scan rate to be calculated.

The approach in the nominal three-gyro system was carried out in the telemetry processing and filtering task (see Chapter 14) as follows:

(1) star mapper data for both the B_T and the V_T channels (sampled at 600 Hz) were extracted from the telemetry frames and placed in a circular buffer;

(2) each data sample from the V_T channel was tested against a threshold (nominally corresponding to V_T = 9 mag). This threshold was adaptive to reflect changes in the background noise level;

(3) a window (nominally 500 data samples) was created for further processing;

(4) the V_T channel data (which had a better signal-to-noise ratio than the B_T channel) was correlated with a four-peaked matching filter function corresponding to an average single-slit response based on the nominal vertical and chevron slit responses (see Figure 5.3). For the star mapper the slit separation in data samples was 20, 60, 40. For

computational efficiency a 'folded' signal was first calculated:

$$E_i = S_i + S_{i+20} + S_{i+80} + S_{i+120} - 4B \qquad (i = 8, N - 127) \qquad [15.10]$$

where S_i is the raw star mapper data sample, B is the background noise, and N is the number of samples in the window. This folded signal was then correlated with a single peaked matching filter for one slit, to provide the full correlation;

(5) the correlation peak was calculated by parabolic interpolation and checks were performed to ensure that sufficient side lobe peaks were found at the correct locations, with an appropriate relationship between the main peak height and the side lobe peak height;

(6) the B_T channel data were correlated using the estimated scan-rate to check that a correlation peak was found with the same datation as for the V_T channel.

The algorithm to determine the scan rate was based on the above method. Essentially steps 4 and 5 were repeated assuming different scan rates, though a modified correlation algorithm was used as defined below. For scan rates where the algorithm showed a star transit, parabolic interpolation of the peak heights as a function of scan-rate was used to give a new maximum peak height; the corresponding scan-rate became the estimate.

The slit system had been defined such that at the nominal scan-rate of 168.75 arcsec s^{-1} a star would cross the slits in 0.2 s. Since the sampling rate was 600 Hz this corresponded to 120 data samples. A scan-rate step of 168.75/120 = 1.40625 arcsec s^{-1} was chosen. This meant that the number of data samples between the first and fourth slits was always an exact integer value. For calculating the scan-rate interpolation problems were avoided by only using data from the outer slits.

For the scan-rate estimation:

$$G_i = S_i + S_{i+120+k} - 2B \qquad [15.11]$$

where:

$$k = \frac{(168.75 - v)}{168.75} \times 120 \qquad [15.12]$$

and v is the variable scan-rate incremented by 1.40625 arcsec s^{-1}.

The output correlated signal was then calculated as:

$$M_i = \sum_{j=-6}^{6} F_j G_{i+j} \qquad [15.13]$$

for all i, where F_j is the response for the average slit derived from the responses for a vertical slit and a chevron slit. This was represented by a function with 13 discrete points. M_i is the output signal which could be used to calculate the scan-rate. To verify the crossing was a star transit and not a noise spike it was necessary to check for side lobes. This could only be achieved if the complete signal was correlated i.e. for the second and third slits as well.

The folded signal for the second and third slits was:

$$H_i = w_1 S_{i+k_1} + w_2 S_{i+k_1+1} + w_3 S_{i+k_2} + w_4 S_{i+k_2+1} - 2B \qquad [15.14]$$

where w_i are the weighting factors dependent on the scan rate being processed, calculated using linear interpolation. k_1 and k_2 are also scan-rate dependent centred around 20 and 80 respectively. Then:

$$N_i = \sum_{j=-6}^{6} F_j H_{i+j} \qquad [15.15]$$

and the sum $O_i = M_i + N_i$ was checked for the requisite number of side lobes. The heights of the side lobes were checked as before. The positions of the side lobes, which are a function of velocity, were also checked.

The estimated scan-rate was updated if there was only one good transit in the data window. This avoided incorrect estimates due to transit coupling or interference between main peaks and side lobes of the other peaks.

As presented above, the scan-rate should have been controllable within ±30 arcsec s^{-1} of the nominal 168.75 arcsec s^{-1}. 51 trial scan-rates were processed by the algorithm. The accuracy of the scan-rate estimation required at this stage was of the order of 0.5 arcsec s^{-1}. This allowed the nominal scan-rate to be achieved with sufficient accuracy.

Once a consistent estimate of the scan-rate had been achieved, the number of scan-rates processed was reduced to 11. This gave a band of approximately 7 arcsec s^{-1} around the latest estimate. This very significantly reduced the processing time required. More side lobes were checked for than in the nominal telemetry pre-processing and star mapper filtering to eliminate spurious harmonics.

Once 5–10 min of closely correlated scan-rate estimates had been achieved, the thruster on time required to correct the scan-rate would be calculated. Before the scan-rate correction was sent, the number of scan rates processed would be increased to cover the nominal scan-rate case. The software would then automatically track the thruster firing and calculate new scan-rates for each transit received. Once the scan-rate had settled down the number of scan-rates could be reduced back to 11.

When a scan-rate close to the nominal value had been reached the software had to calculate an accurate estimate of the full spacecraft attitude. The telemetry pre-processing and star mapper filtering task had now to calculate precisely the same information as for the nominal mission. Thus for each transit an accurate estimate of the star transit time over the virtual slit centre and the B_T and V_T magnitudes were required. Since tasks further down the chain process the z-gyroscope data this was now replaced by the scan rate estimate.

The B_T and V_T magnitudes were calculated using only data from the outer slits, and so avoided inaccuracies due to data interpolation.

Tests were performed at non-nominal scan-rates using the attitude and payload simulator in ESOC, which was adapted for the purpose. Various scan-rates were used. The maximum deviations from the nominal rate used were 30 arcsec s^{-1}, corresponding to the expected maximum error in correcting the scan-rate after perigee. Some results are summarised in Table 15.1.

Scan-rates are measured in arcsec s^{-1}. Over runs of typically an hour or more, star transits of varying brightness taken from the Hipparcos Input Catalogue, were simulated at these different scan-rates. For each star, a rate estimate was produced. Their mean

Table 15.1. Results from on-ground simulations (see text).

Simulated	Mean Error	Standard Deviation
200.0	-0.14	0.30
168.75	-0.09	0.14
138.75	-0.13	0.27

and standard deviation were computed. As expected, the error in determining the correct scan rate grew as the scan rate moved further away from the nominal. However, the accuracy was still good enough even at highly irregular scan-rates to allow for the rate to be corrected by direct thruster commanding.

For the later stages, with near-nominal scan-rates the on-ground Kalman filter was capable of smoothing out irregularities in the rate estimates.

Ground Real-Time Attitude Determination

The remaining tasks in the ground real-time attitude determination chain (described more fully in Chapter 14) are now summarised. One new task was developed to support the ground real-time attitude determination: 'star pattern offset matching' which is described below.

Slit distinction and gyro correction: The slit distinction algorithm remained unaltered from the nominal case. The gyro correction algorithm 'corrected' the transit times of vertical slit crossings to remove the effects of spin-rate variation i.e. it produced a set of pseudo vertical transit times as if the spacecraft had been spinning at a constant rate. Although the algorithm for gyro-correction remained unaltered for the non z-gyro case the (less accurate) scan-rate estimations were used rather than the measured gyro rates; this had implications for the next task in the chain.

Star pattern monitoring: Star pattern monitoring used star pattern recognition to identify each transit pair and determine the field of view of observation. The algorithm mapped the transit times to angles around a great circle. This mapping was less precise when no gyro data was available for the transit time correction. Consequently larger tolerances were needed in the pattern recognition, which raised the probability of confusion or misidentification. An enhanced algorithm was implemented to improve the performance of this task, both in the initialisation phase and during the normal real-time execution. In addition, the user interface was re-worked to provide greater flexibility in obtaining the initial star through parameter options and through improved alphanumeric and graphical displays.

Fine attitude estimation: In the nominal ground real-time attitude determination, a nine element state vector defining Tait-Bryan attitude error angles, body rates and gyro drifts was used in a Kalman filter, receiving measurements from star mapper transit times and the gyros rates at those transit times. Changes were necessary for the new version to allow the star mapper scan-rate estimates to replace those from the z-gyro. Obviously, the estimation of the z-gyro drift was suppressed and the gyro misalignment projection matrix adjusted accordingly. The user interface was altered to allow the spacecraft operator to monitor closely the estimated scan-rate. In tests it was found that a misidentification could have serious impacts for the scan-rate estimate and hence after parameter tuning, a much smoother scan-rate estimate was obtained, obviously at

the cost of sensitivity. Another significant enhancement was the inclusion of actuation on-times to improve ground body rate estimates.

Real-time attitude initialisation: This task provided an instantaneous estimate of the on-board real-time attitude determination state-vector which was valid at precisely the time that the data were accepted on-board.

The real time attitude initialisation software was updated to calculate the values for the new state vector implemented in the on-board real-time attitude determination for the non z-gyro case.

For state vector initialisation:

(i) ϕ_n, θ_n and ψ_n the angular errors, were derived as the difference between the attitude as calculated by fine attitude estimation and the expected attitude from the nominal scanning law;

(ii) transverse gyro rates and drifts were calculated by fine attitude estimation;

(iii) ψ_d was derived as follows:

$$\psi_d = A_d \cos 3\Omega_d t + B_d \sin 3\Omega_d t \qquad [15.16]$$

where Ω_d = 168.56 arcsec s^{-1} (the scan-rate including precession, A_d and B_d were calculated by orbital oscillator routines (see below);

(iv) ψ_2 was set to zero;

(v) then $\psi_1 = \psi_n - \psi_d$;

(vi) $\dot{\psi}$ was provided by fine attitude estimation (from the scan-rate estimate calculated by telemetry pre-processing and star mapper filtering);

(vii) ψ_2 was set to zero;

(viii) $\ddot{\psi}$ was set to zero.

On-Board Torque Model Support

The ground software for computing orbital oscillator parameters refreshed the on-board real-time attitude determination orbital oscillator model, ($\cos\Omega_d t$, $\sin\Omega_d t$) at 15 min intervals (see Chapter 8). For the new on-board real-time attitude determination the deterministic ψ model coefficients, A_d, B_d had also to be maintained. But:

$$\ddot{\psi}_d = \frac{T_z}{I_{zz}} \sin(3\Omega_d t + \gamma_{z0}) \qquad [15.17]$$

where T_z is the torque around the z-axis, I_{zz} is the z inertia, and γ_{z0} the phase of the torque model about z. By double differentiation of the equation for ψ_d, expressions for A_d and B_d were derived. T_z and γ_{z0} were in any case calculated by the ground software.

Star Pattern Offset Matching

This program provided a more efficient way of determining the initial heliotropic angle Ω_0 to within $0\overset{\circ}{.}1$ accuracy. During three-gyro operations, this task was performed in two steps (Chapter 14): field of view separation, and star pattern matching. These two tasks were performed off-line once sufficient star transits had been identified as a result of star mapper filtering and slit distinction. Typically around 20 identified stars were required before a reliable star pattern match could be performed. The process involved considerable manual intervention and a stable scan rate (taking approximately 30 min). Fortunately, during three-gyro operations, Ω_0 remained constant except during rare contingencies. Subsequently, this task was rarely required.

It was clear that, for two-gyro operations, changes of several degrees in Ω_0 could be expected after every perigee. It therefore became productive to look for ways to improve the speed and reliability of this part of the ground procedure. This was achieved by taking a more statistical approach to the determination of Ω_0. The ground system assumed knowledge of the spacecraft z-axis orientation (even if Ω_0 was unknown). Consequently, matching of star mapper transits could be limited to a strip on the sky rather than the complete celestial sphere. The 'star pattern offset matching' algorithm would make no *a priori* assumptions concerning the field of view in which a particular star transit was seen; rather the algorithm would identify all stars around the sky-strip with similar magnitude (in both B_T and V_T) and calculate, for each candidate star, and for both fields of view, the value of Ω_0 which would allow the specified star to be observed in the specified field of view at the measured transit time.

In this way, all the measured star transits gave rise to several candidate values for Ω_0. However when these were binned into small intervals (typically 1°) and displayed as a histogram, a clear peak around the true value was invariably apparent (unless the scan rate was not well controlled to the nominal value at the time).

The real-time 'star pattern monitoring' task was adapted to allow the input value of Ω_0 to be changed on-line. Consequently, the operator at the start of initialisation procedures could attempt to establish a match with 'star pattern monitoring' based on the value of Ω_0 in the previous orbit. Simultaneously the same star data were being used off-line to build up a histogram of potential Ω_0 values. Typically a trend was visible with as few as 10 transits. Assuming the operator was unsuccessful in obtaining a good match from the previous value of Ω_0, a new value could be used in real-time as soon as it became apparent.

Additional Changes

Additional ground software changes to support the two-gyro real-time attitude determination operations were:

(i) modifications to nominal scanning law routines to allow slow precession from sun-pointing to 43° maintaining real-time attitude determination convergence in the event that no other way could be found to recover from an emergency sun pointing manoeuvre taking place. This modification was never used operationally;

(ii) an additional command was developed to reset the deterministic z-axis error angle on-board. This allowed the inhibition of the z-axis precession around the 43° cone for use during long losses of signal. This operation is described more fully below;

(iii) gyro configuration checking was suppressed in other gyro monitoring software;

(iv) programme star file observation covariance matrices used for the modulation strategy (see Chapter 8) were re-calculated for the extended mission;

(v) given that several attitude and orbit control system software patches were now active, a new facility was developed for the ground system to perform the selection and automatic uplink of the software patches to the spacecraft. This was first used for the two-gyro patch and the decoupling of the z-thrusters.

15.5. Operational Experience

Parallel Tests During Three-Gyro Operations

The on-board software was designed to run in parallel with the nominal real-time attitude determination software. It was implemented and had been initialised on several occasions, for performance testing, using the new ground real-time attitude determination system. Figure 15.2 shows a graph of the nominal and new ψ Tait-Bryan error angles over 40 min on 26 September 1991.

The nominal real-time attitude determination was converged throughout this period (as in all tests performed with the spacecraft). However, at the start of this graph, the new real-time attitude determination was diverged. After 13 min, ground real-time attitude determination software had calculated and reset the new state vector to the correct value. For the remainder of this period, the two methods kept good agreement, even across two z-thruster firings (shown by the discontinuous changes in gradient).

It was also interesting to compare the performance of the individual components making up the new ψ estimate over the same period. Figure 15.3 shows the values of ψ_1 and ψ_d. ψ_d is the sinusoidal curve. The effectiveness of the ψ_d model in describing the disturbance torques in the apogee region was seen by the fact that the ψ_1 term was nearly linear between thruster firings. ψ_2 remained much smaller throughout. In tests, the new real-time attitude determination was found to remain converged throughout the apogee region (i.e. for several hours). Like the nominal real-time attitude determination, it was however incapable of surviving the perigee region, below 6000 km, where the disturbance torques were greatest.

It was necessary to close the shutters on the payload detectors, whenever the Earth or Moon crossed one of the fields of view. During these occultations, real-time attitude determination ran on gyro and thruster data alone. With three operational gyros, this presented no problems. However, without a z-gyro, z-thruster on-times affected the rate estimates after a firing. Figure 15.4 shows one such firing taking place during an occultation with a subsequent divergence of the new real-time attitude determination from the nominal.

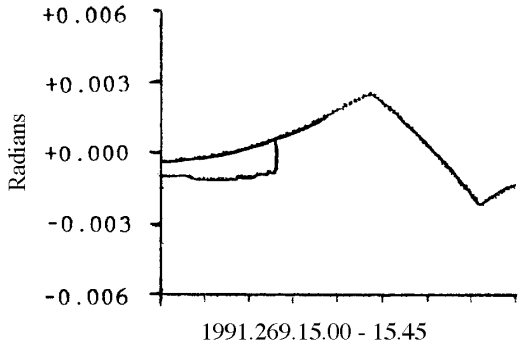

Figure 15.2. *Convergence of the z-axis Tait-Bryan error angle ψ after ground real-time attitude determination correction.*

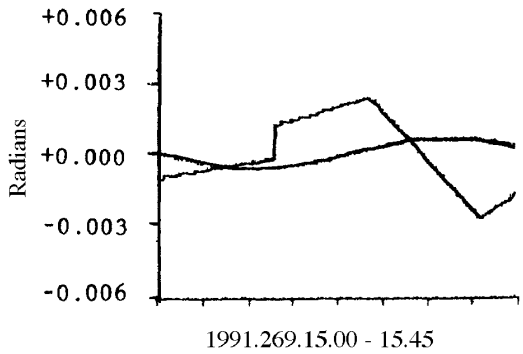

Figure 15.3. *Contribution of ψ_1 (linear part) and ψ_d (smooth curve) to the total Tait-Bryan error angle ψ over the same period as in Figure 15.2. ψ_d is the periodic term reflecting the evolving disturbance torques. ψ_1 is the linear term which was corrected from ground, and which changed gradient after each thruster actuation.*

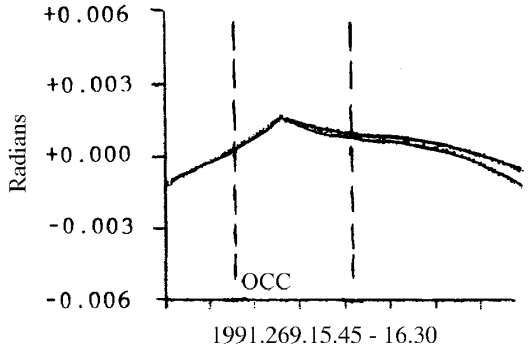

Figure 15.4. *Divergence of ψ after a thruster firing during an occultation (the region within the vertical dashed lines, OCC), where no star transits were available to correct the rate estimates.*

Full Two-Gyro Operations

Once two-gyro operations resumed after the hibernation period, it was seen that the performance of real-time attitude determination, when converged, was greatly improved, as measured by the size of the innovations (once more generally below 1 arcsec). This could be directly attributed to the substitution of a noisy gyro 5, by a smooth, although simplified, disturbance torque model around the spacecraft z-axis. This torque model

was designed specifically to model the radiation pressure torque at apogee. As a result, real-time attitude determination performance would deteriorate as the altitude of the spacecraft dropped near perigee due to the increase in size in un-modelled torques and high star mapper background noise. During perigee, there was no possibility that real-time attitude determination would remain converged due to a total absence of information about orientation around the z-axis: no gyro, a blind star mapper and an uncalibrated disturbance torque model. It was expected that the spacecraft would spin-down or spin-up through perigee by as much as 30 arcsec s^{-1} every orbit.

Hibernation studies had suggested that it might be possible to partially predict in an empirical manner the likely amount of spin-up/down. In practice, this proved to be ineffective. Initially, much higher changes in spin rate were seen (50 – 60 arcsec s^{-1}). This was traced in part to futile attempts by the on-board normal mode control software to take erroneous real-time attitude determination output and fire the z-thrusters, sometimes exacerbating the change in spin rate. A procedure was put in place, involving resetting various real-time attitude determination variables on-board which completely disabled the z-thruster around perigee. This decreased the size of the spin-rate changes, as well as saving gas. However, predictions still proved unreliable and manual compensation through perigee for possible spin rate changes was abandoned. Correction of the spin rate was earlier made easy by the use of gyro 4. Although too unreliable for use with real-time attitude determination, the deteriorated gyro output could be used when telemetry was re-acquired after perigee, while the star mapper was still blinded and unusable by the ground software.

On 22 December 1992, gyro 4 failed completely, thereby making it necessary to run ground real-time attitude determination to obtain spin-rate estimates based on the four peak separation of a star mapper transit. Although the on-ground star mapper filter was capable of determining the scan rate over a wide range of potential values, it was inefficient in the high noise regions of van Allen belts. Moreover, occasional deviations from the nominal scan rate of up to 50 arcsec s^{-1} were still apparent making an initial estimate very difficult. Extra facilities were quickly added to allow a graphical display of particularly bright star mapper transits. This allowed a crude visual estimate of the scan rate to be obtained, even in high noise, so that scan rate corrections could be achieved as quickly as possible after perigee.

At the start of two-gyro operations, some difficulty was found in initialising the on-board real-time attitude determination system from ground. The errors in the on-ground attitude determination, particularly in the rate estimation around the z-axis, were higher than for three-gyro operations. Given that the delay of between 30 and 60 s from the time of last good star observation to the execution time of the on-board initialisation command, errors were building in the on-ground extrapolation of the on-board state vector. The result was that the attitude was sufficiently accurate to be within the 30 arcsec innovation threshold on-board, but produced a large initial innovation before bringing the on-board estimate within the narrow threshold range of 10 arcsec. This innovation was also used, to update the transverse gyro drifts (see Equation 15.8) when in fact the drifts were reasonably stable over the short period of the initialisation. The now erroneous drift values subsequently led to erratic convergence of the on-board system or even divergence again. Work continued to improve the accuracy of the ground software. However an additional measure was taken by devising a second set of fixed gains for use on-board. This set was used during initialisation procedures and contained much smaller gains for updating the gyro drifts. The first gain sets produced for the two-gyro operations had assumed gyro drift gains 100 times higher

than those used during three-gyro operations. Values closer to the previous sets used in three-gyro operations were re-adopted for the new 'extended window' set. These safeguarded the gyro drifts during initial convergence. Once convergence was achieved and the innovation threshold was set to narrow range, a more sensitive set of gains was implemented on-board to produce smoother performance of the attitude determination during science data collection.

Further tuning of the gains was carried out as the performance of the remaining gyros worsened to lessen the effects of erratic drifts diverging the on-board system. In addition calibration studies by ESOC and the Royal Greenwich Observatory (RGO) resulted in more accurate z-thruster performance figures being used on-board.

After 1 November 1992, it was necessary to make occasional changes to the initial precession angle defining the z-axis position on the 43° cone around the sun vector. This was to safeguard the spacecraft through long losses of signal (more than 5 hours around perigee). During a long loss of signal, the nominal scanning law precession rate was set to zero. Without this, a spin-up/down of the spacecraft through perigee could have resulted in a large change to the initial heliotropic angle, which defined the orientation of the spacecraft x and y axes. In the worst case, the initial heliotropic angle could have changed by up to 90°. In this case the last orbital oscillator command sent before loss of signal would have driven the z-axis perpendicular to the nominal scanning law. Left without correction over a typical lost pass, (12 hour loss of signal), the z-axis could drift in solar aspect angle by more than 3° and would consequently have come dangerously close to triggering an emergency sun reacquisition order.

Figure 14.15, showing the overall science data return, clearly indicates that significantly less data were collected per orbit. Overall 'good' data were returned 37 per cent of the total time, or 51 per cent of the covered time. There were three main reasons for this: (1) the real-time attitude determination initialisation procedure was required every orbit; (2) despite the improvement made during the hibernation period, the initialisation procedure was still lengthier than during three-gyro operations; (3) real-time attitude determination divergence after occultations was rather frequent (over 50 per cent of the time), usually occurring over the longer ones where a thruster firing took place in the middle.

Occultations occurred in seasons: sometimes only at perigee, thereby securing uninterrupted coverage for most of the orbit; sometimes occurring in pairs every revolution of the spacecraft, making data return exceptionally poor. Analysis was performed by the NDAC Consortium and ESOC into the possibility of slewing the spacecraft around the 43° nominal scanning law cone to avoid occultations, but it was concluded that this would be detrimental to the overall observation programme.

Software changes were made to improve the performance of the ground real-time attitude determination software immediately after occultations in an effort to reduce the time to re-initialise real-time attitude determination. Steps were also taken to understand the disturbance torques and thruster performance better in the hope of improving the probability of real-time attitude determination survival through occultations. Such improvements were still on-going when further gyro failures necessitated the use of the zero-gyro real-time attitude determination system.