14. REAL-TIME ATTITUDE DETERMINATION WITH THREE GYROS

This chapter describes the details of the attitude determination and control using the nominal three-gyro configuration. In acquiring the scanning law—a procedure which would have been required just once in the nominal geostationary scenario, but which had to be repeated many hundreds of times throughout the revised mission (attitude knowledge was lost during perigee passages and during periods of high star mapper background levels)—a series of procedures had to be followed: initial star transit recognition, star pattern matching, and precise attitude determination. Since the procedures had to be used repeatedly, considerable effort went into their optimisation.

14.1. Introduction and Overall Concept

The spacecraft was constrained to rotate once every 2 hours 8 min about its principal spin axis, and this spin axis precessed about a 43° half-cone centred on the Sun. This motion was called the nominal scanning law, and was defined in a heliotropic reference frame (see Chapter 8 for details). The heliotropic attitude angle Ω was defined with respect to an initial heliotropic angle Ω_0, which was not known before the initial star pattern recognition. The actual orientation of the satellite (x, y, z frame) with respect to the predefined nominal scanning law (X, Y, Z frame) was described using the Tait-Bryan error angles ϕ, θ, ψ shown in Figure 14.1.

In practice, the spacecraft was allowed to deviate from the nominal scanning law by up to 500 arcsec, before being corrected by the cold-gas thrusters (approximately every 10 min). This ‘normal mode’ control is described in more detail in Chapter 13.

Since the image dissector tube could only sample one star at a time, the instantaneous field of view was controlled by the on-board computer to track the programme stars for short intervals, switching rapidly from one to the other. The computer had to be able to calculate the expected position of a star on the grid at any given time. Excessive memory would have been required to hold the entire catalogue on-board permanently and to calculate the expected times. Instead a memory buffer of transit times (the programme star file, described in detail in Chapter 8) for the next 30 min of observations was maintained and regularly refreshed by ground command. The transit times were

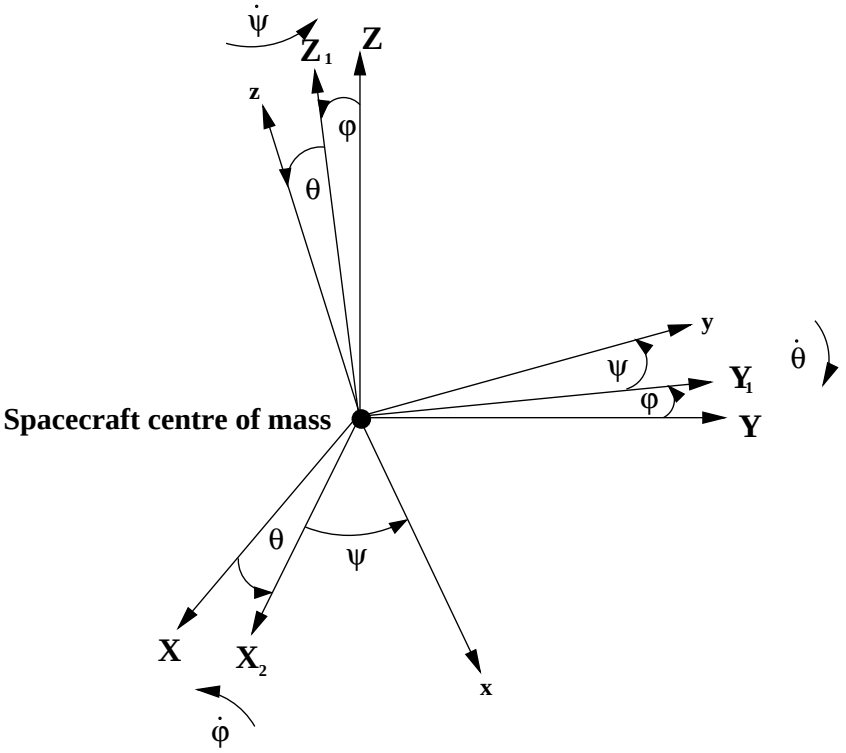

Figure 14.1. *Tait-Bryan error angles of the spacecraft frame (x, y, z) relative to the nominal scanning law frame (X, Y, Z).*

calculated in advance on ground at ESOC, assuming that the spacecraft was perfectly following the nominal scanning law.

The on-board computer therefore required detailed knowledge of the attitude, both for spacecraft control and for image dissector tube piloting. This was achieved through on-board software called the real-time attitude determination system, which received inputs from three operational rate-integrating gyros every 1.066 . . . s. The gyro data, after pre-processing, allowed the determination of the rotation rates about the three principal orthogonal axes of the spacecraft. Real-time attitude determination was able to integrate these measurements, taking into account knowledge of the thruster firing times to predict the attitude at any time. This alone however did not provide sufficient accuracy for the image dissector tube piloting. The gyros themselves were subject to biases (known as gyro drifts) in their rate estimation which led to a steadily increasing error in the attitude determination. The real-time attitude determination was capable of compensating for these gyro drifts assuming they were accurately determined. In practice, these drifts varied with time due to thermal distortion from the external environment and 'random walk' effects. The only way for real-time attitude determination to calibrate the drifts was to use a second independent data source.

This source came from the measurements of the stars across the star mapper grid, consisting of two sets of four aperiodically spaced slits (see Chapter 2). One set was vertical with respect to the scanning direction, the other arranged in a ±45° chevron (see Figure 2.10). The errors between the expected and measured transit times of the stars across these two sets of slits were used, via Kalman filter techniques on-board, to refine the attitude determination to approximately 1 arcsec accuracy. The modulated

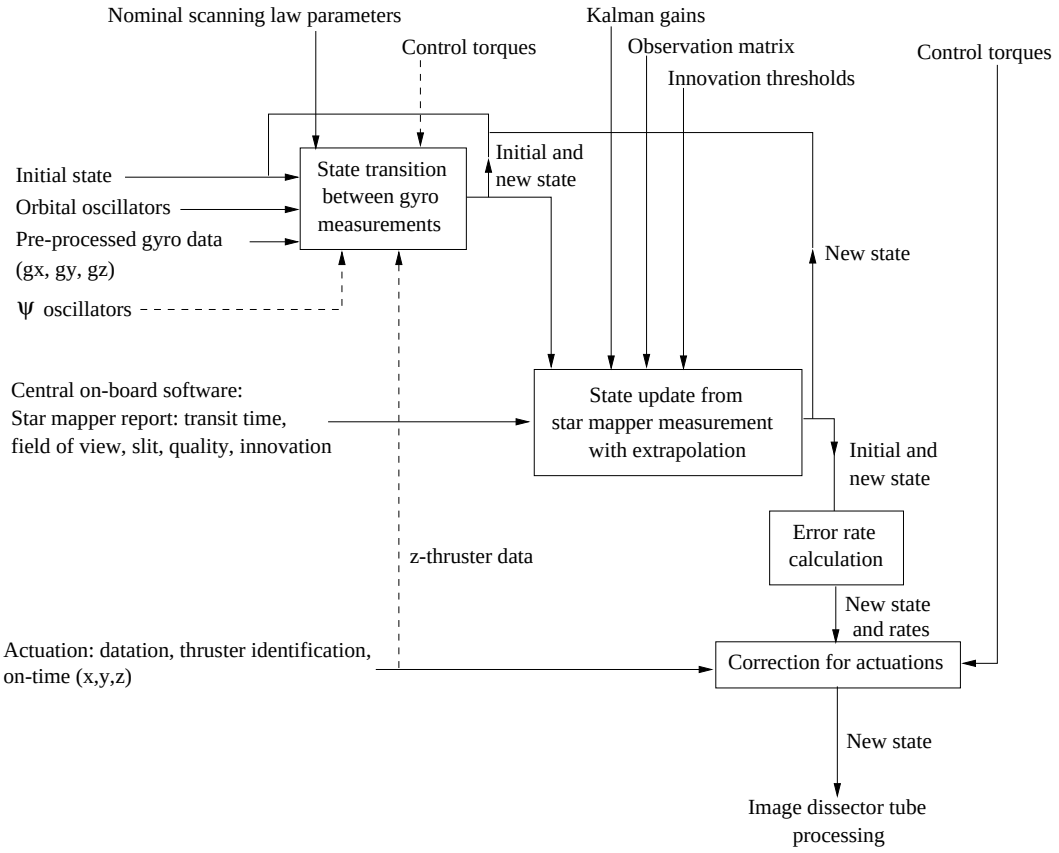

Figure 14.2. *The interconnection of on-board tasks related to the real-time attitude determination using gyros. The dashed lines indicate the additional input required for the implementation of the two-gyro attitude control described in Chapter 15.*

light from the star mapper grid was sampled by two photomultiplier tubes operating in two different spectral bands.

Each set of four slits was separated in a 1:3:2 ratio. Thus when starlight crossed the four slits, a distinctive pattern was discerned by the on-board software, which filtered the raw data assuming an ideal single-slit response profile. The resulting modulated signal was characterised by one main lobe and several side lobes. The height and width of these lobes were checked against various accuracy criteria before being passed to real-time attitude determination for use in determining the attitude. Such tests included determining the height of the main lobe against the value expected for a star of that brightness.

The various elements of real-time attitude determination are shown in Figure 1.5, and the interconnection of on-board tasks related to the real-time attitude determination using gyros is illustrated in Figure 14.2.

When a star had been successfully identified in one or other of the fields of view, the transit time was measured by the position of the main lobe in the filtered data. By comparison with the expected transit time as given in the programme star file, a measurement of the attitude error was derived using a Kalman filter. In this filter, under three-gyro control, six values were estimated: three angles which defined the error

between the current spacecraft attitude and the nominal scanning law (the Tait-Bryan error angles), and three gyro drifts converted into rates along the principal spacecraft axes.

As a result of tuning since the start of the mission involving ESOC, Matra Marconi Space and the scientific consortia, the real-time attitude determination system was capable of working to an accuracy of 1 to 2 arcsec in attitude determination if stars were being successfully identified in the star mapper processing. However, if errors had built up to over 10 arcsec, star measurements could no longer be used by real-time attitude determination, which diverged from the real attitude. This happened after 50 per cent of all perigee passes, where higher disturbance torques and noise in the star mapper signal (caused by van Allen belt radiation) made divergence more likely.

As a protection against falsely identified stars, the attitude determination on board normally rejected any star transit times which showed that the spacecraft attitude relative to the nominal scanning law had drifted by more than 10 arcsec since the last star transit was processed by real-time attitude determination. During normal operations this was not expected to happen, because the gyro drifts were accurate to less than 5 milliarcsec s^{-1} and real-time attitude determination received star transit data every 10 s on average.

Should anomolous conditions cause such an error to build up, it was possible to increase this 10 arcsec limit ('narrow window') to 30 arcsec ('extended window') to allow real-time attitude determination to use these larger transit time errors to update its attitude knowledge. If the errors were greater than 30 arcsec, then intervention was required from ground, because the on-board software was no longer capable of processing the correct buffer of raw star mapper data for stars to be detected. A suite of programs called the ground real-time attitude determination system was developed and used at ESOC for this purpose. This software was designed to process on-ground all star mapper data and estimate the attitude independently of any real-time attitude determination results.

14.2. On-board Real-Time Attitude Determination using Three Gyros

On-board real-time attitude determination involved estimating in real-time a state vector consisting of:

- ϕ: Tait-Bryan error angle around the spacecraft x-axis;
- θ: Tait-Bryan error angle around the spacecraft y-axis;
- ψ: Tait-Bryan error angle around the spacecraft z-axis;
- d_x: gyro drift around the spacecraft x-axis;
- d_y: gyro drift around the spacecraft y-axis;
- d_z: gyro drift around the spacecraft z-axis.

Data from three gyros gave angular velocities (spacecraft body rates) about all three axes. This allowed a complete prediction of the attitude from one processing cycle (16/15 s) to another, assuming the previous attitude estimate was correct. Small corrections to these predictions were made by comparing expected and measured star mapper crossing times.

Real-Time Attitude Determination Definitions

The following variables were input to the real-time attitude determination algorithm:

$G_x(k+1)$, $G_y(k+1)$, $G_z(k+1)$: angular measurements currently available from the gyro data cyclic acquisition software.

t_d: star detection time derived from the on-board star mapper processing.

R: star angular measurement derived from the on-board star mapper processing.

(ϵ, f): star mapper slit identifier of the detected star:
- $\epsilon = 0$: vertical slit
- $\epsilon = -1$: inclined slit, upper half of field of view
- $\epsilon = +1$: inclined slit, lower half of field of view
- $f = +1$: preceding field of view
- $f = -1$: following field of view

IQ: star quality:
- IQ = 1 for a quality 1 star (more accurate on-ground position, typically better than 0.4 arcsec)
- IQ = 2 for a quality 2 star (less accurate on-ground position, typically in the range 0.4–0.8 arcsec)

T: time of real-time attitude determination estimate (multiple of the real-time attitude determination cycle time t_i).

$S_v(k)$, $C_v(k)$: sine and cosine values produced by the orbital oscillator, but corresponding to the previous time interval (see Chapter 8).

$t_{on,x}$, $t_{on,y}$, $t_{on,z}$: the X, Y, and Z cold-gas thruster on-times provided by the nominal mode control algorithm (see Chapter 13).

Ω_T: transverse rate defined by Equation 8.10.

Ω_N: estimate of the absolute Z-axis nominal scanning law spin rate provided by the on-board computer.

The following on-board constants were used:

t_i: the cycle time (16/15 s) between inputs from the inertial reference unit.

KG_1 (6,6): the 6×6 matrix of fixed gains associated with quality 1 stars.

KG_2 (6,6): the 6×6 matrix of fixed gains associated with quality 2 stars.

One column of KG_1 or KG_2 was used in a star mapper update of the real-time attitude determination filter estimate. The column used was determined by the slit identifier as follows:
- column 1: $(\epsilon, f) = (0, 1)$
- column 2: $(\epsilon, f) = (-1, 1)$
- column 3: $(\epsilon, f) = (1, 1)$

– column 4: $(\epsilon, f) = (0, -1)$
– column 5: $(\epsilon, f) = (-1, -1)$
– column 6: $(\epsilon, f) = (1, -1)$

The rows corresponded to the six components of the filter state vector:

$$V = \begin{pmatrix} \phi(k+1) \\ \theta(k+1) \\ \psi(k+1) \\ d_x(k+1) \\ d_y(k+1) \\ d_z(k+1) \end{pmatrix} \qquad [14.1]$$

where $d_x(k+1)$, $d_y(k+1)$, $d_z(k+1)$ were the current estimates of the gyro drifts.

H (6,6): star mapper observation matrix. The row used in a star mapper update was determined by the slit identifier (exactly as for the column of KG_1). The columns corresponded to the six components of the filter state.

T_{ax}, T_{ay}, T_{az}: control torques applied about the spacecraft x, y and z axes.

I_{xx}, I_{yy}, I_{zz}: principal moments of inertia about the x, y and z axes.

μ: innovation threshold for narrow and extended windows.

Defining additionally:

$\omega_X(k)$, $\omega_Y(k)$, $\omega_Z(k)$: angular rates with respect to inertial space expressed in scanning law X, Y and Z axes and relating to the previous time interval.

$d_x(k+1)$, $d_y(k+1)$, $d_z(k+1)$: current estimates of drift rates about the spacecraft x, y, and z axes.

R_I: the innovation to be used in the filter star update, i.e. the difference between the extrapolated star mapper measurement and the linear approximation of the star mapper measurement using the present real-time attitude determination attitude estimates.

The following outputs were obtained:

$\phi(k+1)$, $\theta(k+1)$, $\psi(k+1)$: current estimates of attitude error angles.

$\dot{\phi}(k+1)$, $\dot{\theta}(k+1)$, $\dot{\psi}(k+1)$: current estimates of attitude error rates.

$\phi_m(k+1)$, $\theta_m(k+1)$, $\psi_m(k+1)$: modified (predicted) estimates of attitude error angles around actuation.

$\dot{\phi}_m(k+1)$, $\dot{\theta}_m(k+1)$, $\dot{\psi}_m(k+1)$: modified (predicted) estimates of attitude error rates around actuation.

Three-Gyro Real-Time Attitude Determination Algorithm

At each sampling interval the real-time attitude determination produced estimates of the Tait-Bryan angles and their derivatives from the output of the gyro data cyclic acquisition software (see Figure 14.2).

The first step was the calculation of the nominal scanning law rates about the theoretical X and Y axes using the output of the orbital oscillator:

$$\begin{aligned} \omega_X(k) &= \Omega_T \; S_v(k) \\ \omega_Y(k) &= \Omega_T \; C_v(k) \\ \omega_Z(k) &= \Omega_N \end{aligned} \qquad [14.2]$$

where $\Omega_T = 168^{\circ}.75 \text{ hr}^{-1}$, and $\Omega_T = \Omega_p \sin\xi$ with $\xi = 43^{\circ}$ (see Chapter 8 on the orbital oscillator and its scheduling). The estimates of the Tait-Bryan angles were then calculated according to the matrix estimate equation:

$$\begin{pmatrix} \phi(k+1) \\ \theta(k+1) \\ \psi(k+1) \\ d_x(k+1) \\ d_y(k+1) \\ d_z(k+1) \end{pmatrix} = \begin{pmatrix} 1 & \omega_z t_i & 0 & -t_i & 0 & 0 \\ -\omega_z t_i & 1 & 0 & 0 & -t_i & 0 \\ 0 & 0 & 1 & 0 & 0 & -t_i \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \phi(k) \\ \theta(k) \\ \psi(k) \\ d_x(k) \\ d_y(k) \\ d_z(k) \end{pmatrix}$$
$$+ \begin{pmatrix} G_x(k+1) - \omega_X(k)\; t_i - G_y(k+1)\; \psi(k) \\ G_y(k+1) - \omega_Y(k)\; t_i + G_x(k+1)\; \psi(k) \\ G_z(k+1) - \omega_Z(k)\; t_i - G_x(k+1)\; \theta(k) \\ 0 \\ 0 \\ 0 \end{pmatrix} \qquad [14.3]$$

Finally, estimates of the first derivatives of the Tait-Bryan angles were calculated by:

$$\begin{aligned} \dot{\phi}(k+1) &= (\phi(k+1) - \phi(k))/t_i \\ \dot{\theta}(k+1) &= (\theta(k+1) - \theta(k))/t_i \\ \dot{\psi}(k+1) &= (\psi(k+1) - \psi(k))/t_i \end{aligned} \qquad [14.4]$$

When star information was available from the on-board computer it was used to update the gyro-based estimates of the Tait-Bryan angles and their derivatives. As a result of the necessary processing time, there was a delay between the detection time and the time the star information was available in the real-time attitude determination. Therefore, the star measurement was extrapolated to make it consistent with the time of the most recent gyro measurement. This extrapolation was a function of the change in $\phi(k+1)$, $\theta(k+1)$ and $\psi(k+1)$ over the relevant time interval:

$$\begin{aligned} R_E = R &+ H(*,1)\; \dot{\phi}(k+1)\; (T - t_d) \\ &+ H(*,2)\; \dot{\theta}(k+1)\; (T - t_d) \\ &+ H(*,3)\; \dot{\psi}(k+1)\; (T - t_d) \end{aligned} \qquad [14.5]$$

where the $*$ represents the relevant row of the observation matrix as determined by the slit identifier. If the extrapolation procedure had been invalidated by a thruster actuation, the star information was ignored. When this was not the case the difference between the extrapolated star measurement and its linear approximation was given by:

$$R_I = R_E - (H(*,1)\; \phi(k+1) + H(*,2)\; \theta(k+1) + H(*,3)\; \psi(k+1)) \qquad [14.6]$$

If this innovation was less than the innovation threshold μ, the Tait-Bryan angles together with the estimates of the gyro drift were updated according to the filtering equation:

$$\begin{aligned}
\phi(k+1) &= \phi(k+1) + \mathrm{KG}_1(1,*)\ R_I \\
\theta(k+1) &= \theta(k+1) + \mathrm{KG}_1(2,*)\ R_I \\
\psi(k+1) &= \psi(k+1) + \mathrm{KG}_1(3,*)\ R_I \\
d_x(k+1) &= d_x(k+1) + \mathrm{KG}_1(4,*)\ R_I \\
d_y(k+1) &= d_y(k+1) + \mathrm{KG}_1(5,*)\ R_I \\
d_z(k+1) &= d_z(k+1) + \mathrm{KG}_1(6,*)\ R_I
\end{aligned} \qquad [14.7]$$

If the innovation was greater than the innovation threshold the star information was ignored. In Equation 14.7 the update is associated with a quality 1 star. For a quality 2 star KG_1 was replaced by KG_2 in the above equations.

Additional Attitude Prediction in the Case of Actuation

The output of the real-time attitude determination was used in the attitude and orbit control system control loop and in the pointing of the image dissector tube beam. As the normal pointing of the image dissector tube beam used information two time cycles out of date, an additional attitude prediction was necessary around a thruster actuation to prevent the situation where the image dissector tube beam was pointed without accounting for the thruster firing.

The real-time attitude determination outputs corresponding to the observation frame $(k+1)$ preceding the actuation, and to the observation frame $(k+2)$ including the actuation were extrapolated from the nominal real-time attitude determination estimates according to equations of the form:

$$\begin{aligned}
\dot{\phi}_m(k+2) &= \dot{\phi}_m(k+1) = f(\dot{\phi}(k+1),\ t_{\mathrm{on},x},\ T_{ax},\ I_{xx}) \\
\dot{\theta}_m(k+2) &= \dot{\theta}_m(k+1) = f(\dot{\theta}(k+1),\ t_{\mathrm{on},y},\ T_{ay},\ I_{yy}) \\
\dot{\psi}_m(k+2) &= \dot{\psi}_m(k+1) = f(\dot{\psi}(k+1),\ t_{\mathrm{on},z},\ T_{az},\ I_{zz})
\end{aligned} \qquad [14.8]$$

and:

$$\begin{aligned}
\phi_m(k+2) &= \phi_m(k+1) = g(\phi(k+1),\ \dot{\phi}(k+1),\ \dot{\phi}_m(k+1),\ t_{\mathrm{on},x},\ t_i,\ T_{ax},\ I_{xx}) \\
\theta_m(k+2) &= \theta_m(k+1) = g(\theta(k+1),\ \dot{\theta}(k+1),\ \dot{\theta}_m(k+1),\ t_{\mathrm{on},y},\ t_i,\ T_{ay},\ I_{yy}) \\
\psi_m(k+2) &= \psi_m(k+1) = g(\psi(k+1),\ \dot{\psi}(k+1),\ \dot{\psi}_m(k+1),\ t_{\mathrm{on},z},\ t_i,\ T_{az},\ I_{zz})
\end{aligned} \qquad [14.9]$$

Modified Attitude Prediction

If during a given cycle k, the scanning law control determined that an actuation was necessary this actuation was synchronised with the observation frame cycle $(k+1)$. The real-time attitude determination outputs corresponding to cycle k preceding the actuation and to cycle $(k+1)$ including the actuation were extrapolated from real-time attitude determination estimates and from applied thruster on-times. The real-time attitude determination outputs corresponding to cycle $(k+2)$ following the actuation were then nominal again.

For cycle k, preceding the actuation, real-time attitude determination outputs were given by:

$$\begin{aligned}
\dot{\phi}_m(k) &= \dot{\phi}(k) + T_{ax} t_{\mathrm{on},x} \\
\dot{\theta}_m(k) &= \dot{\theta}(k) + T_{ay} t_{\mathrm{on},y} \\
\dot{\psi}_m(k) &= \dot{\psi}(k) + T_{az} t_{\mathrm{on},z} \\
\phi_m(k) &= \phi(k) - \tfrac{1}{2} T_{ax} t_{\mathrm{on},x}^2 - T_{ax} t_{\mathrm{on},x} t_i \\
\theta_m(k) &= \theta(k) - \tfrac{1}{2} T_{ay} t_{\mathrm{on},y}^2 - T_{ay} t_{\mathrm{on},y} t_i \\
\psi_m(k) &= \psi(k) - \tfrac{1}{2} T_{az} t_{\mathrm{on},z}^2 - T_{az} t_{\mathrm{on},z} t_i
\end{aligned} \qquad [14.10]$$

For cycle $(k+1)$, including the actuation, real-time attitude determination outputs were given by:

$$\begin{aligned}
\dot{\phi}_m(k+1) &= \dot{\phi}(k+1) + T_{ax} t_{\mathrm{on},x} \\
\dot{\theta}_m(k+1) &= \dot{\theta}(k+1) + T_{ay} t_{\mathrm{on},y} \\
\dot{\psi}_m(k+1) &= \dot{\psi}(k+1) + T_{az} t_{\mathrm{on},z} \\
\phi_m(k+1) &= \phi(k+1) - \tfrac{1}{2} T_{ax} t_{\mathrm{on},x}^2 \\
\theta_m(k+1) &= \theta(k+1) - \tfrac{1}{2} T_{ay} t_{\mathrm{on},y}^2 \\
\psi_m(k+1) &= \psi(k+1) - \tfrac{1}{2} T_{az} t_{\mathrm{on},z}^2
\end{aligned} \qquad [14.11]$$

The modified estimates were computed only for on-board computer acquisition purposes. They were never used in the nominal attitude prediction.

14.3. On-Ground Real-Time Attitude Determination

The initial on-ground attitude initialisation could be split into two distinct phases:

Initial Star Pattern Recognition: This phase generated off-line pairings between observed star transits and star identifications from a catalogue, identifying both the slit system crossed and the field of view of observation. This allowed the initial attitude of the scanning law to be determined, and provided initialisation for the next phase of the attitude determination.

The on-ground initial star pattern recognition was divided into five stages, each of which passed its output data as input to the next stage. These stages were:

(a) telemetry pre-processing and star mapper filtering;
(b) slit distinction and gyro correction;
(c) field of view separation;
(d) star pattern matching against a star catalogue;
(e) initialisation of the nominal scanning law.

All of the above steps were necessary if Ω_0 was unknown (as it was at the start of the mission). Generally however, under three-gyro control, attitude drifts after perigee were less than 100 arcsec. Therefore in general, ESOC maintained a constant value of Ω_0 throughout the routine (three-gyro) phase of the mission. The only times when it was deemed necessary to change Ω_0 during this period were during a small number of contingencies where the spacecraft assumed a non-nominal spin rate for a significant period of time. In order to re-acquire attitude control, ESOC had the choice of options, either to perform a large firing of the cold gas thrusters to slew the spacecraft around its

principal rotation axis thereby re-acquiring the previous nominal scanning law attitude, or to calibrate the value of Ω_0 for the current phase and calculate a new programme star file. In all cases where the drift around the rotation axis was more than one or two degrees, the second option was faster to implement, safer and less expensive in fuel. This implied that in those cases where Ω_0 was known, only the first two stages up to 'slit distinction and gyro correction' needed to be taken before moving on to precise attitude determination.

Precise Attitude Determination: This activity paired, in close to real-time, observed star transits and star identifications from the operational catalogue, identifying the slit system crossed, whether the upper or lower half of the chevron slit system was crossed, and the field of view of observation. This information was used to generate a precise attitude in near real-time and to estimate the gyro drifts. Linear extrapolation was used to predict the spacecraft attitude at a precise time (approximately 25 s in the future) and the attitude along with the gyro drifts were uplinked to the spacecraft to be received on-board at precisely the time at which the attitude was valid. This uplink allowed initialisation of the on-board real-time attitude determination.

The precise attitude determination was divided into four distinct stages, each of which passed its output data as input to the next stage. These stages were:

(a) star pattern monitoring;
(b) fine attitude estimation;
(c) basic angle calibration;
(d) initialisation of the on-board real-time attitude determination.

The basic angle calibration was a subsidiary task, only used at the start of the mission and thereafter infrequently. The other stages were all critical to real-time attitude determination initialisation.

Telemetry Pre-processing and Star Mapper Filtering

This stage of the initial star pattern recognition involved the collection of both star mapper and attitude and orbit control system telemetry data, together with processing of the star mapper data. This was a time-critical task which ran throughout the attitude determination. The raw data from the satellite telemetry were received as telemetry 'bundles'. Each bundle contained 0.666... s of raw telemetry data equivalent to 16 telemetry formats (1 format = 41.666... ms). Data from the on-board attitude and orbit control system required by tasks in the attitude initialisation chain were extracted and made available to all tasks.

The star mapper samples for both the B_T and the V_T channels were also extracted and placed in a circular buffer. Every raw star mapper sample in the V_T channel was tested against a threshold (usually corresponding to $V_T = 9$ mag). This threshold was adaptive to reflect changes in the background noise level, for example as experienced as the spacecraft passed through the van Allan radiation belts. A sample exceeding this threshold could be from a peak from a star crossing one of the slits (vertical or inclined). Therefore a window of data around this sample was correlated with a four-peaked matching filter function which corresponded approximately to the expected signal from a bright star. The uneven spacing of each slit system (and also of the matching filter function) gave one correlation peak with distinct determinable side lobe peaks for a valid star transit.

A transit was identified if a correlation peak was found with the same time datation in both the B_T and the V_T channels. The height and location of the peaks allowed a transit time and approximate B_T and V_T magnitudes to be determined, which were passed to the next task in the initialisation chain. It was of course possible to have peaks from more than one transit after correlation. In this case each transit was identified by the algorithm except where there was interference between main peaks and side lobes of other peaks.

Figure 14.3 shows the star mapper data for the V_T channel contained in a 0.666... s bundle of raw telemetry data. The crossing by a bright star of one of the slit systems is clearly seen and would have been detected by the correlation process. Figure 14.4 also shows the star mapper data for the V_T channel contained in a 0.666... s bundle of raw telemetry data. The solitary large spike would have exceeded the raw threshold thus triggering the star mapper filtering. However in this case no peaks would have been detected after correlation. Figure 14.5 shows similar data for the B_T channel in which multiple transits can be seen.

Slit Distinction and Gyro Correction

The transit data from the filtering process described above did not contain entirely accurate information, since the correlation function was not optimal for a given transit crossing for the following reasons: (i) the exact focus of the telescopes was not known at the start of the mission; (ii) the filter was matched to the nominal rotation rate of the satellite, so that deviations from the nominal rate would smear the response; (iii) the star mapper photometric performance and slit characteristics were different for the vertical slits and the inclined slits. There was no *a priori* information about the slit system crossed and hence a correlation function had to be used for an 'average' slit performance.

The first two error sources led to a magnitude bias while the third led to a bias between the magnitude calculated by a star transit of the vertical slit system and the magnitude calculated by a transit of the same star through the inclined slit system. Successive transits of the same star image across the vertical and inclined slit systems were separated by between 0.3 and 7.2 s (see Figure 2.10). A star crossing the centre of the star mapper slit system would transit the second slit systems 7.2 s after it transited the first whereas this time separation was 0.3 s for one crossing at the extremes of the star mapper. The slit distinction algorithm attempted to 'pair' transits separated by no more than 7.2 s by comparison of their magnitudes. Two transits occurring within 7.2 s of each other were paired if their B_T and V_T magnitudes and colour ($B_T - V_T$) were equal within certain tolerances; these tolerances being set to take account of the errors referred to above. Each member of the pair was then tested separately to ensure that it could not form part of any other pair. Any pairs containing a transit that could belong to more than one pair were discarded and their data were not used by following tasks.

This process led to approximately 40 per cent of transits being discarded. Real-time slit distinction was not possible: on receipt of data from a bright star transit crossing the first slit system at most 7.2 s delay could occur before data were received from the second slit crossing. Besides, following 'pairing' of two transits a further delay of 7.2 s after reception of the second transit of the pair was needed to ensure that the pairing was unique. Once a successful pairing had been found a revised magnitude estimate could be formed by taking the arithmetic mean of the two magnitudes of the two transits

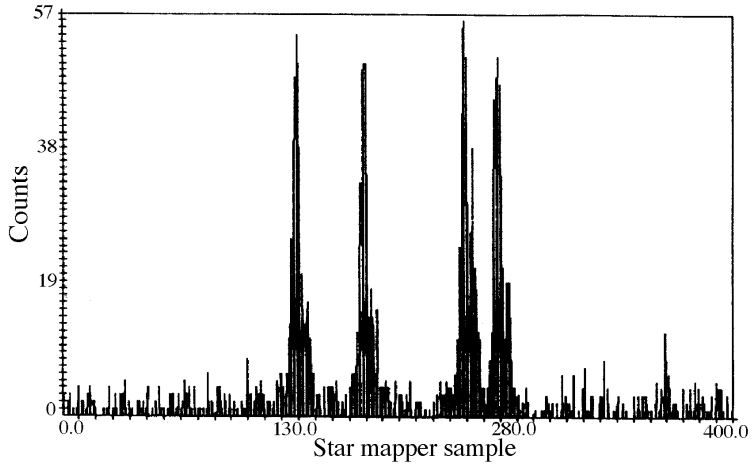

Figure 14.3. *A typical bright star crossing the four slits of the star mapper, as sampled in the V_T channel in a 0.666 . . . s 'bundle' of raw telemetry.*

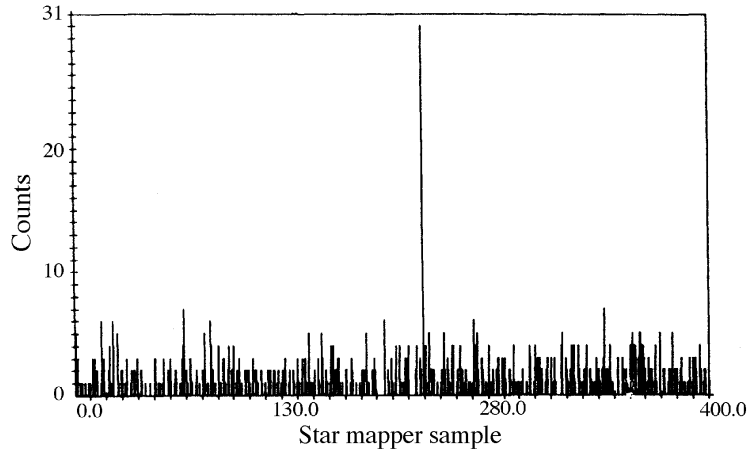

Figure 14.4. *Star mapper data (V_T channel) showing the occasional appearance of solitary noise spikes in the star mapper telemetry. Such an outlier would trigger star mapper filtering, but would not result in a false star transit being detected.*

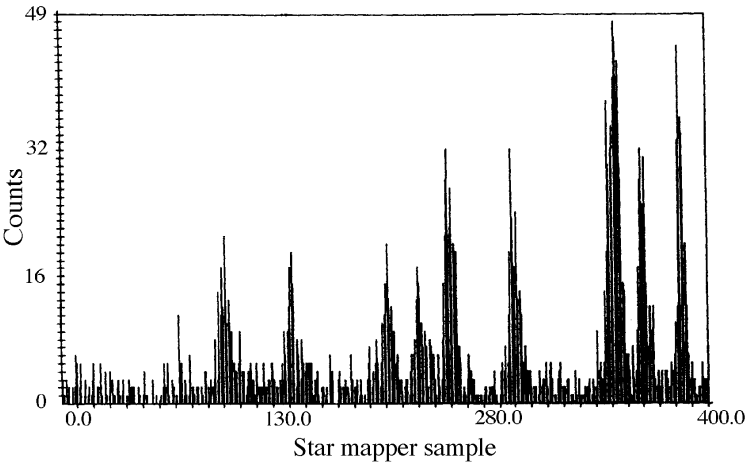

Figure 14.5. *A typical example of several stars crossing the different star mapper slits (B_T channel) at nearly the same time. The ground real-time attitude determination system could, in almost all cases, separate the multiple star transits without confusion.*

forming the pair. As later tasks wanted to match transits between fields of view and with a star catalogue a map from a vertical crossing to a phase round a great circle was performed. This map could be improved if variations in the spacecraft scan rate as detected by the on-board z-gyro were used to give a corrected vertical slit crossing time. This correction was performed at this stage by accessing the attitude and orbital control system gyro data stored by the telemetry pre-processing.

Field of View Separation

As explained above, this stage was only necessary if Ω_0 was completely unknown.

The two fields of view of the payload scanned approximately the same strip of sky with a phase difference equal to the basic angle. In practice the following field of view did not exactly scan the same strip of the celestial sphere as the preceding field of view, although there was generally sufficient overlap to match patterns of stars observed in the two fields of view.

Data from the slit distinction algorithm were collected over a time interval of around 1 hour. The data thus collected were used to construct two fields of equal duration whose start times were separated by the minimum time over which the basic angle could be scanned—about 20 min. These were called field 1 and field 2.

A succession of time offsets were added to the gyro corrected transit times in field 1. These time offsets were such that: (i) each offset was within a predefined tolerance around the nominal time taken to rotate through the basic angle; (ii) after adding the time offset to the transit times in field 1 at least one transit in field 1 had exactly the same time associated with it as a transit time in field 2. All such time offsets were considered.

For each offset a count was made of the number of transits that 'aligned', i.e. whose transit times, magnitude and colour were equal within specific tolerances. For alignment counts above a predefined threshold of 5 a residual was calculated which was equal to the root of the mean sum of squares of the time separations between aligned transits.

For all time offsets which produced an alignment count greater than 70 per cent of the maximum count, that which resulted in the lowest residual was selected as the best offset and the transits from field 2 which were matched were taken as an identified pattern of stars from the following field of view only.

In practice the field of view separation algorithm proved itself to be both robust and accurate, even though it was not essential to have error-free results at that stage. As a by-product of the field of view separation the average scan rate over the basic angle was calculated.

Star Pattern Matching Against a Star Catalogue

As explained above, this stage was only necessary if Ω_0 was completely unknown.

The transits identified by the field of view separation algorithm as being in the following field of view were used as input to the star pattern matching algorithm. This task attempted to identify these transits by matching with stars from a 360° catalogue strip. The pole of this strip was given by the predicted z-axis of the satellite, corresponding to

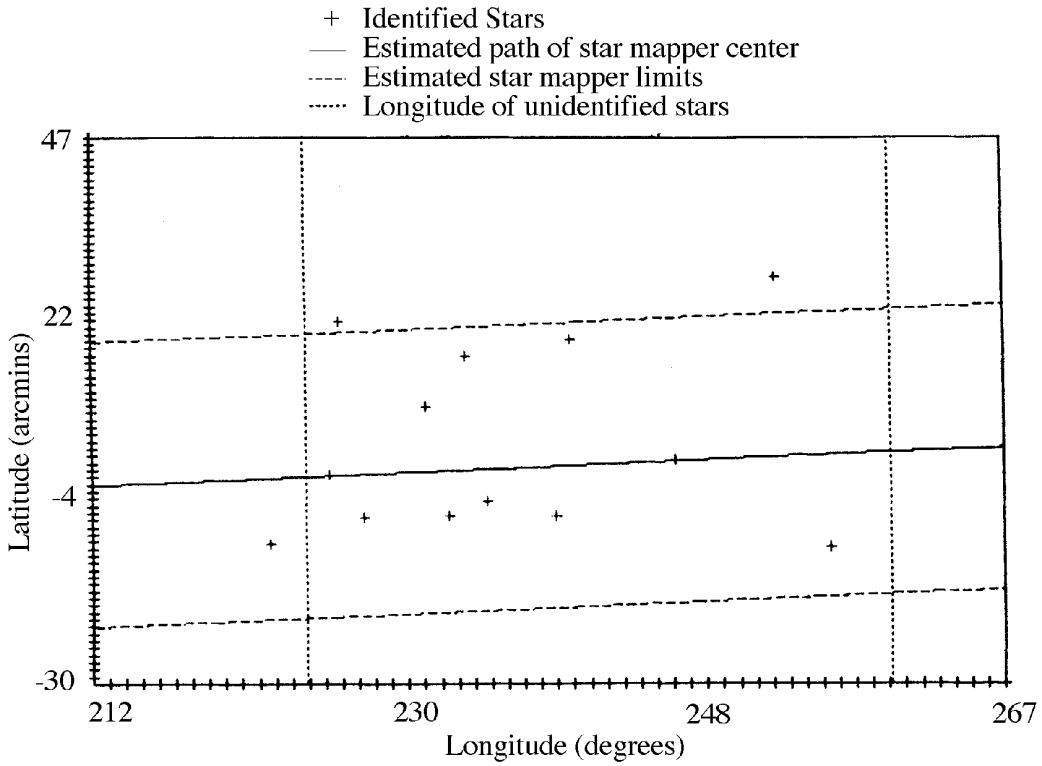

Figure 14.6. *An example of graphical output from the ground real-time attitude determination star pattern matching facility. Star transits were paired, where possible, with stars in the star catalogue strip (see text for details). Longitude and latitude were measured with respect to the pole of the star catalogue strip. Successfully identified stars are denoted by crosses. Unidentified stars are denoted by vertical 'longitude' lines.*

the nominal scanning law Z-axis. It was found in practice that a maximum strip width of 2° was needed for the algorithm.

The gyro corrected vertical transit times were mapped onto a great circle using the rotation rate estimate derived from the field of view separation results.

A similar technique to that used in field of view separation was then applied to match the pattern of star transits with a star catalogue strip covering the possible scanned band of the celestial sphere. To avoid possible confusion within the algorithm stars of similar magnitude separated by less than 2 arcsec along the strip were removed before matching was attempted. All valid offset angles between 0 and 360° were considered and the best offset determined using the same criteria as defined for field of view separation.

Because this matching was CPU-intensive the processing was performed on a fast off-line machine, with input data transfer through a high-speed computer link. The algorithm was demonstrated to be accurate in operations, coping with any possible 'missing' stars in both the catalogue and the set of transits. To give the user confidence in the match given by the algorithm the location of the identified stars along the strip were plotted. A 'best fit' path of the star mapper centre was drawn and the 40 arcmin star mapper extent shown. Figure 14.6 shows an example of this. Unidentified transits are depicted by a dashed vertical line.

Initialisation of the Scanning Law

Again, as explained above, this stage was only necessary if Ω_0 was completely unknown.

The value of the Ω_0 of the nominal scanning law was not known before initial star pattern recognition. However if Ω_t is the heliotropic angle derived from the nominal scanning law for the attitude initialisation zero time assuming $\Omega_0 = 0$, Ω_m is the offset calculated from the following field of view to the star catalogue strip, Ω_s is the offset of the centre of the star mapper vertical slit system from the spacecraft *x*-axis, then the initial heliotropic angle is given as: $\Omega_0 = \Omega_m - \Omega_t + \Omega_s$ and thus from this point on the scanning law could be determined.

Star Pattern Monitoring

Whereas the initial star pattern recognition identified detected transits retrospectively, the aim of star pattern monitoring was to identify transits as close to real time as possible. It did this by maintaining a sliding window of six matched vertical transits in each field of view. New transit pairs which had been identified by the slit distinction algorithm in near real-time were received. The vertical transit time was mapped to an angle round the great circle and added to the existing pattern in both fields of view. A pattern match was attempted on each field of view using a similar—though not identical—algorithm to that used to calculate Ω_m, based on a small range of offsets centred around the previous offset for each field of view.

If the transit was identified in one and only one of the fields of view it was accepted. The window for that field of view was updated by releasing the 'oldest' star and adding the new identification. The transit pair data and star identification were passed to the fine attitude estimation task.

Once an accurate value for Ω_0 was known (either from previous spacecraft operations or as a direct result of the full initial star pattern recognition just described) the two field of view windows required by star pattern monitoring were easily initialised. The offset calculated in star pattern matching was used to calculate offsets to the catalogue strip for each field of view. The set of following field of view transits from field 2 of the separation algorithm (see field of view separation above) was used to initialise the following field of view, and the remaining transits in field 2 were assumed to be rich in preceding field of view transits and these were used to initialise the preceding field of view. If the on-board real-time attitude determination needed to be subsequently re-initialised, providing the nominal scanning had been maintained, the star pattern monitoring could be re-initialised using the nominal scanning law to determine the offsets to the catalogue strip for each field of view.

Up to 20 transit pairs identified by slit distinction were used as trial patterns for both fields of view and pattern matching was attempted. If a pattern was identified uniquely in both fields of view then the star pattern matching task was initialised and could begin near real-time identifications. Otherwise further transit pairs were collected from slit distinction and further pattern matches attempted. Figure 14.7 shows the results from initialising one of these windows in this manner. Here a succession of closely spaced offsets, each giving an alignment count of 11 transits, can be seen. The residual is plotted on top of this and the local minimum, corresponding to the best pattern match,

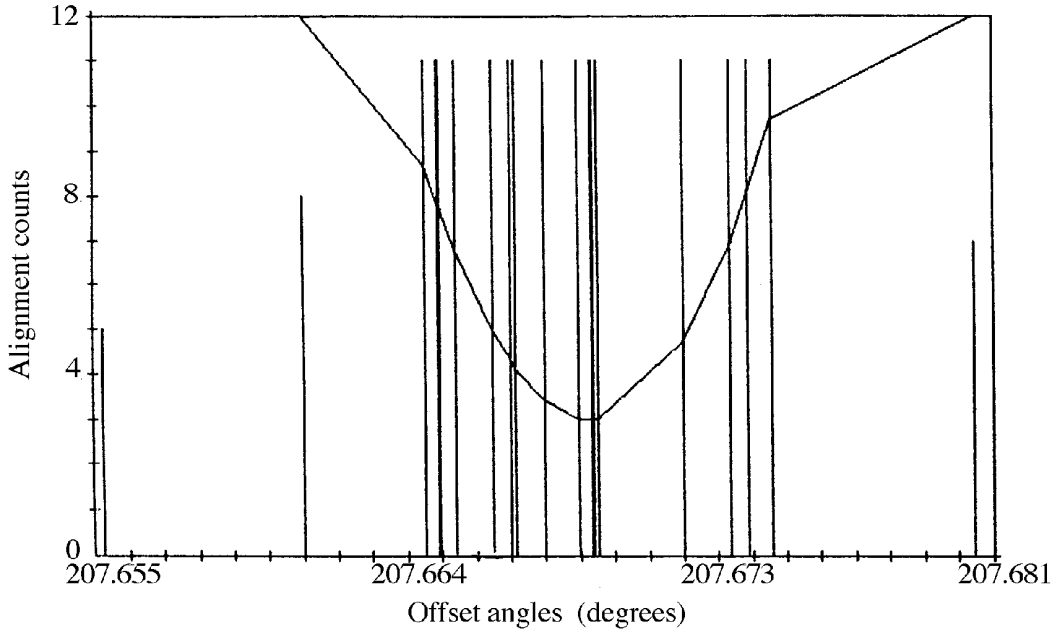

Figure 14.7. *An example of graphical output from the ground real-time attitude determination star pattern monitoring facility, during the initialisation phase. The angular offset from the strip origin which gives the best pattern match is that showing the maximum alignment count and lowest residual (see text for details).*

is located toward the centre of this offset grouping. (This is also the behaviour observed for field of view separation and star pattern matching.)

The star pattern monitoring task ran throughout the remainder of the operations. A graphics display was used to illustrate the real-time identifications. An example is shown in Figure 14.8. The display is split up into 5 sections: the top panel displays the incoming transits as a function of time. Each incoming star transit is represented by a separate point. The symbol used for the star shows whether the transit has been identified and if so in which field of view. The two central graphs display the identifications in each field of view. The left hand graph displays the following field of view and the right hand graph the preceding field of view. Each identified transit is plotted in star catalogue coordinates in the appropriate field of view. The path of each field of view across the catalogue strip can therefore be seen—this was useful for determining the best half-width to use for the strip. The lower two graphs show the residuals for the two fields of view as functions of time. The residuals will be small—any increase in the residuals shows a deterioration in the quality of the pattern match possibly because of star misidentification.

Fine Attitude Estimation

The requirement for the fine attitude estimation algorithm was to estimate in near real time the spacecraft attitude, body rates and gyro drift rates. At this stage, the attitude was determined by defining three Tait-Bryan attitude angles (ϕ, θ, ψ) in an inertial reference frame.

(ϕ, θ) defined the scan plane orientation and ψ defined the rotational phase of the spacecraft. Hence (ϕ, θ) oscillated about zero, while ψ changed secularly at the same

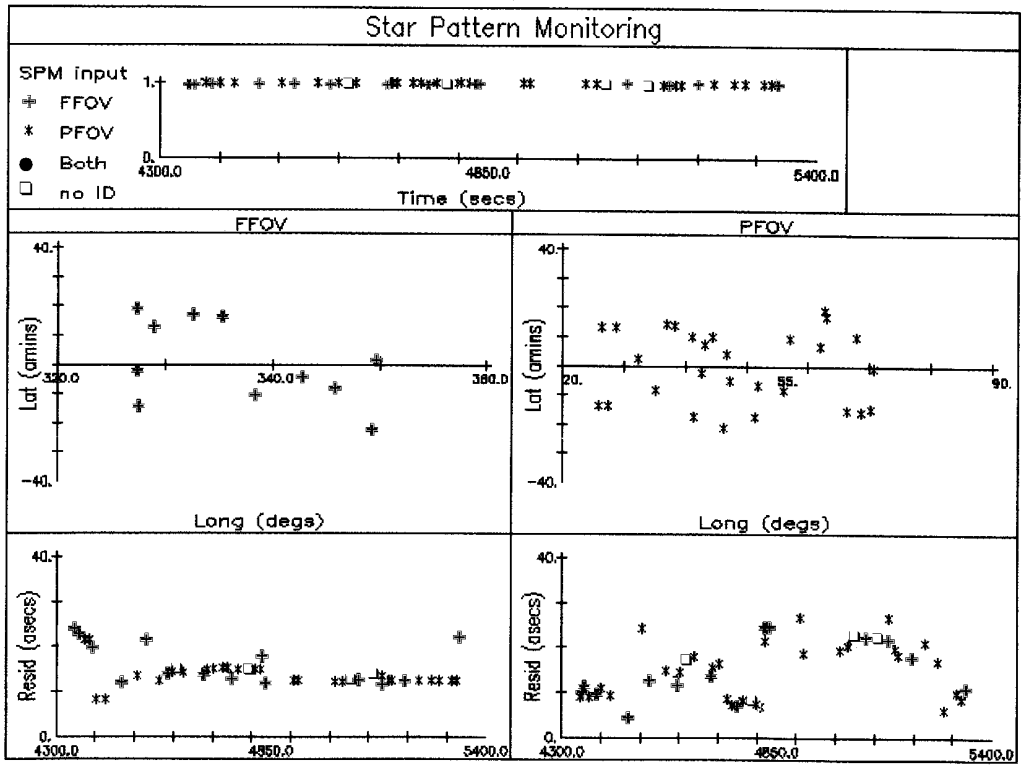

Figure 14.8. *An example of graphical output from the ground real-time attitude determination star pattern monitoring facility, during the routine monitoring phase (see text for details).*

rate as the rotation rate of the spacecraft (168.75 arcsec s^{-1}). The three axis rotation rates were determined from the three operational gyro outputs.

The attitude and rates were determined using the identified star transit pairs (inclined and vertical transits), generated by the star pattern monitoring program, with gyro rate and thruster actuation data extracted from the attitude and orbit control system telemetry stream by the telemetry pre-processor. This information was processed in real-time using a Kalman filter method, to update the spacecraft state estimate, where the state was described by: (i) the attitude (Tait-Bryan angles in either sun-pointing or 43° inertial reference frame); (ii) the body rates about the spacecraft axes, (ω_x, ω_y, ω_z); and (iii) the gyro drift rates about the gyro input axes.

The initial state estimate assumes: $\phi = \theta = 0$; $\psi = 180°$; $\omega_x = \omega_y = 0$; $\omega_z = 168.75$ arcsec s^{-1} and gyro drift rates from pre-calibrated values or previous attempts to initialise real-time attitude determination.

Besides the state estimation processing, the current attitude knowledge was applied to the identified transit pair to determine which inclined slit, upper or lower, was crossed. Transits lying within approximately 5 arcmin of the spacecraft equator were rejected. As the estimation converged, this rejection deadband was reduced, more transits could be processed and the estimation accuracy thereby increased.

The filter processing consisted of the following steps: (a) state and state covariance transition between star transit pairs, with allowance for thruster actuations between

transit pairs; (b) pre-processing of star transit pair data and gyro rates to generate a measurement for input to the Kalman filter; and (c) state vector and covariance matrix update resulting from a new measurement.

State and State Covariance: The state vector and state covariance transition between star mapper measurements was performed as a series of transitions over intervals without actuations, with covariance updates at actuations. Between measurements a discretised linear state model was used, which could be derived from kinematic and dynamical (Euler equations) relationships for the attitude and the body rates respectively. The drift rates were modelled as a random walk with the drift noise incorporated into the model's system noise. Gyro reaction torques and solar radiation pressure torques were taken account of in the forcing function which was used to perform the transition of the state vector. Other disturbance torques (e.g. gravity gradient and magnetic) were modelled as torque noise within the overall system noise.

The performance of the Kalman filter could be tuned by varying the power spectral densities in the system noise matrix. There were six power spectral densities which could be independently tuned: three for the torque drifts and three for the gyro drifts. The power spectral densities used affected the performance of the filter; if they were too large the filter tended to track the raw measurements (including noise), if they were too small the filter response would be slow and the estimates would lag behind the optimal values. The effects of actuations were allowed for by updating the state covariance at the time of the actuation. Specifically, the rate covariance on the axis about which an actuation torque acts was relaxed—the covariance was reset to a comparatively large preset value. The state and the remaining state covariance components were not updated.

Measurement pre-processing: This covered the calculation of a measurement, measurement observation matrix and measurement covariances given a star transit pair and/or the attitude and orbit control system gyro output history. The different measurement types generated were:

(i) transverse measurement (ζ_m), giving the longitude of the star measured from the spacecraft equator, which was a function of:

- inclined/vertical star transit times;
- estimated attitude (ϕ, θ, ψ) in the catalogue strip reference frame at the vertical slit transit time;
- estimated body rates at the vertical slit transit time;
- star latitude and longitude in the catalogue strip reference frame;
- inclined/vertical slit mounting angles in the spacecraft reference frame.

(ii) phase measurement giving the phase of the spacecraft x-axis measured from the catalogue strip reference x-z plane about the scan axis direction. Whereas star transit pairs for both preceding fields of view and following fields of view were processed to give transverse measurements, only transits of one field of view were processed to give phase measurements. This measurement incorporated the scan phase and scan rate knowledge. By considering only one field of view, possible bias effects due to uncertainty in the basic angle separating the two fields of view were removed from the scan phase and scan-rate estimates. After initial calibration of the basic angle at the start of the mission (see next section), this bias was removed and it was safe to use measurements from both fields of view for updating the state vector;

(iii) gyro rates for the two equatorial and the z-axis gyro at a given time.

The measurements were treated as scalars rather than as a 5-dimensional measurement vector.

State Vector and Covariance Matrix Update: The updating of the state and state covariance from the measurement, measurement covariances and observation matrices was performed using the standard Kalman Filter algorithms. The state vector was updated using the differences (innovations) between the observed measurements and those predicted from the previous estimate of the state. The star pattern monitoring may have sent a small number of misidentified stars which led to unexpectedly large differences in the measurements. If these differences had been used the convergence of the filter would have been badly disrupted. To avoid this, a measurement innovation threshold could be set for each of the three types of measurement. These thresholds could be altered during the operation of fine attitude determination to protect the accuracy of the estimates. Using the conventional Kalman Filter equations, it was found that the gyro drift estimates were particularly sensitive to high innovations. Reasonably accurate initial estimates were obtained for these drifts and a way needed to be found to protect those values from excessive alteration at the start of fine attitude estimation when the initial attitude and rates were not so well known. The solution was to modify the measurement innovations going into the drifts estimates by factors which could be tuned during operation of fine attitude estimation. Initially these values were set low to fix the drifts at a constant value until such times as the Tait-Bryan angles and body rates have begun to converge. They could then be increased, once the innovation thresholds were low enough to protect against star misidentifications, so that more accurate final values for the drifts could be obtained.

Convergence of fine attitude estimation was usually achieved within 20 min, but times varied considerably depending on the frequency of star transits. Convergence was improved by having a mixture of transits in both fields of view. Examples of the Tait-Bryan angles, body rates, gyro drifts rates and innovation levels are given in Figure 14.9 to Figure 14.12 for a typical real-time attitude determination initialisation attempt at 43° from sun-pointing.

Basic Angle Calibration

As described in the previous section, a difference between the actual and the measured basic angle (nominally 58°) led to a bias in the predicted attitude (particularly the Tait-Bryan angle ψ). The solution was to collect transits once fine attitude estimation was converged and process them using a least-squares method to derive a better estimate of the basic angle, which could then be used to further improve the convergence of fine attitude estimation. The process consisted of three activities: collecting star transit and attitude information from fine attitude estimation; pre-processing the fine attitude estimation results to obtain a 'batch' of measurements; and estimation of the basic angle and grid rotation parameters by the application of a batch weighted least-squares method.

Once fine attitude estimation was approximately converged, it was possible to collect transit information including: transit time; field of view (following or preceding); catalogue position of observed star; current orientation of the scan plane as estimated by fine attitude estimation; and current rotation rate about the spacecraft z-axis.

These data were collected over 40–60 min by which time approximately 100 measurements should have been collected. During this period only phase measurements from

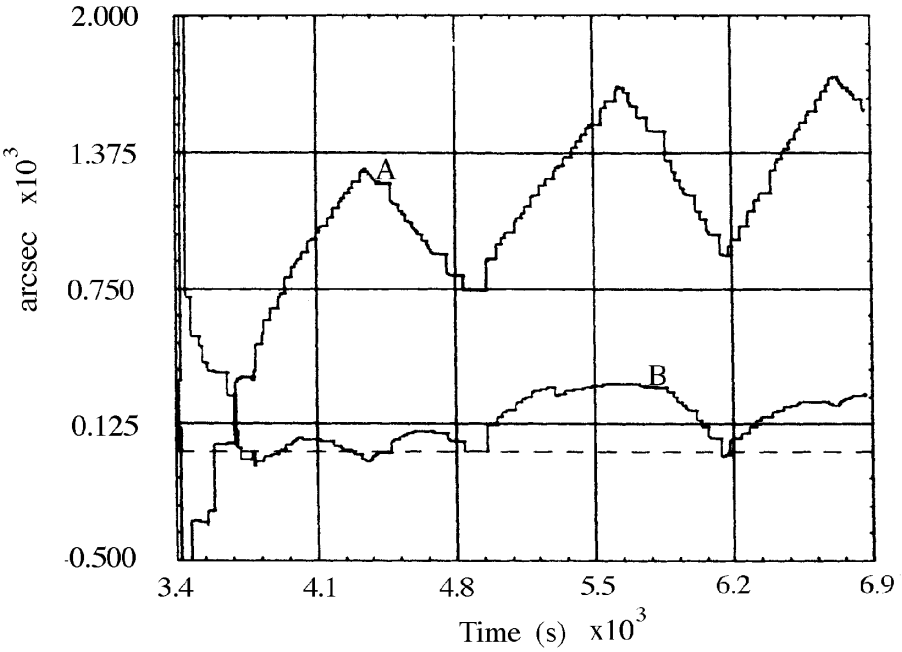

Figure 14.9. *On-ground estimates of the ϕ (A) and θ (B) Tait-Bryan attitude angles from the ground real-time attitude determination fine attitude estimation facility. The discontinuities result from thruster actuations.*

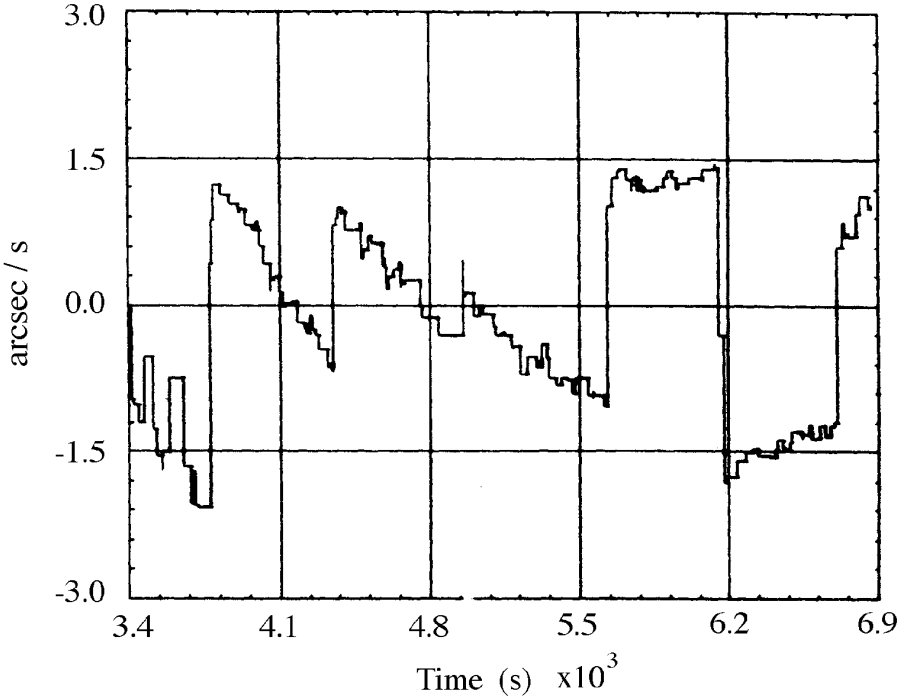

Figure 14.10. *On-ground estimate of body rate about the spacecraft y-axis from the ground real-time attitude determination fine attitude estimation facility.*

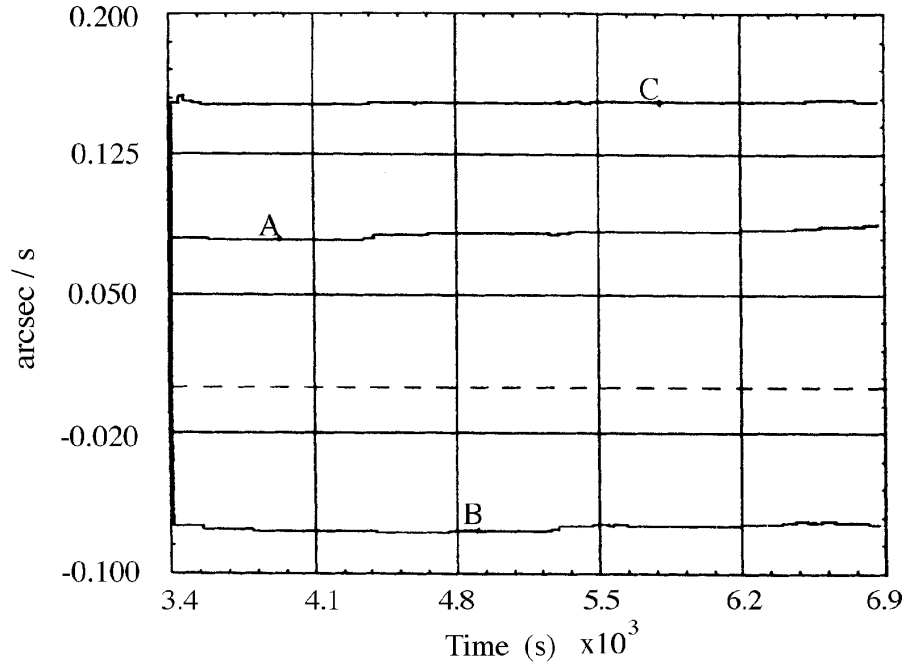

Figure 14.11. *On-ground estimates of the drift rates of gyro 1 (A), 2 (B) and 4 (C), from the ground real-time attitude determination fine attitude estimation facility.*

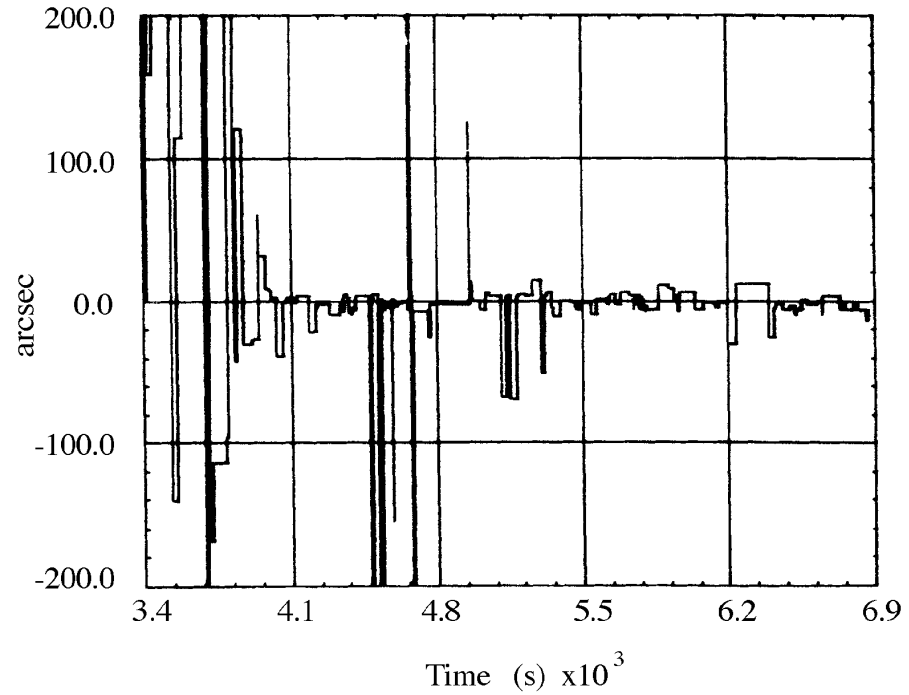

Figure 14.12. *On-ground transverse measurement innovation used in the Kalman filter within the ground real-time attitude determination fine attitude estimation facility. Convergence of the algorithm to a stable attitude estimate is seen as the size of the innovations reduces towards zero.*

one field of view could be used to update the state vector. This eliminated possible biases in the state vector estimates.

At the end of this collection period, the raw measurement data were pre-processed. Near-simultaneous (within a specified time tolerance) transits of the vertical star mapper slits of the two fields of view (preceding and following) were required. For each preceding/following pair selected, the expected star separation was calculated from the known star catalogue positions. The transverse coordinates of each star at the corresponding vertical slit transit time were calculated using the attitude estimated by the fine attitude estimation program and the catalogues star positions. The mean scan rate over the interval spanned by the results stored on disk was computed by integrating the instantaneous z-axis body rate results—as calculated by fine attitude estimation—with time and dividing by the interval duration. These measurements were then combined in a least-squares formulation which provided estimates of the rotation angle of the grid (nominally zero) and the basic angle.

The basic angle was calibrated at the start of the mission while the spacecraft was sun-pointing. The value obtained was 58° 30 arcsec, compared to an on-ground calibrated value of 58° 42 arcsec. The grid rotation estimate proved to be unreliable. The maximum error of 3 arcmin was in any case shown in simulations not to impede convergence of fine attitude estimation. Convergence of fine attitude estimation during operations assumed that the value was zero and no problems had been met. The new value for the basic angle was thereafter used by fine attitude estimation. Subsequent attempts to re-initialise real-time attitude determination were performed using phase measurements from both fields of view without impeding convergence. Within a few weeks of the start of the mission, the data reduction consortia had derived a more accurate value for the basic angle, which was used by the ground real-time attitude determination system.

Initialisation of the On-Board Real-Time Attitude Determination

The final stage of the on-ground attitude initialisation was to provide for the spacecraft an instantaneous attitude estimate in terms of the Tait-Bryan error angles plus the three gyro drifts projected onto the spacecraft axes.

The spacecraft attitude estimated on-ground at any given time would have been derived from data collected on-board the spacecraft at some earlier time. Although this time delay could be calculated by using the computer clock and the correlation between the spacecraft clock and UTC, further time delays would occur in uplinking the attitude to the spacecraft. The following solution was adopted.

A future time, t_1 was selected according to the following criteria: t_1 was at least 20 s in the future; it corresponded exactly to an on-board telemetry format start time (1 telemetry format = 10.666... s); and the total delay from the time, t_0, of on-board acquisition of the data used for the attitude estimate to the time t_1 was calculated. The attitude estimate and the estimated body rates calculated by the fine attitude estimation task and valid for time t_0 were used to linearly extrapolate through the time interval (t_1 - t_0) to give an attitude estimate valid for time t_1. The attitude estimate for time t_1 was converted to the ecliptic reference frame and compared with the nominal scanning law attitude for time t_1 to derive the Tait-Bryan error angles ϕ, θ, ψ. These Tait-Bryan error angles together with the gyro drift estimates calculated by fine attitude estimation were uplinked to the spacecraft.

The uplink command was time-tagged to expire 5 s before the time t_1, but was format synchronised so that the command was only executed at the start of a format. This ensured that the uplinked attitude estimate was accepted on-board at precisely the time at which the estimate was valid.

Note that the above method did not remain valid if there was an actuation of the on-board thrusters during the interval $t_1 - t_0$, as this would invalidate the extrapolation. Thruster actuations were monitored on-ground and the uplink was only performed during a period when no actuations were expected. The Tait-Bryan error estimate had to be within 28 arcsec of the true attitude for the on-board real-time attitude determination to converge following the uplink. To meet this requirement it was essential to keep the time interval $t_1 - t_0$ as small as possible. This illustrates the time-criticality of the on-ground processing from the telemetry pre-processing and star mapper filtering task through to the final attitude estimation. The value of $t_1 - t_0$ was monitored during operations and the uplink was only permitted when this value was less than 60 s. Under conditions of normal computer load this criterion was easily met. Figure 14.13 and Figure 14.14 show the effect of the attitude initialisation uplink on the on-board real-time attitude determination.

Subsequent Ground Operations

Following the uplink the on-board real-time attitude determination took over control of the spacecraft. During the period of attitude initialisation a programme star file was prepared on ground containing expected star transit data and this was uplinked to the spacecraft before the attitude initialisation. If the value of Ω_0 was already known, the programme star file could already be generated in advance of the start of real-time attitude determination initialisation procedures. In general, therefore, during the three-gyro operations, the programme star file was generated automatically every 24 hours on the off-line mainframe computer, transferred to the real-time system and uplinked to the spacecraft, regardless of the status of real-time attitude determination.

The on-board real-time attitude determination was enabled after the attitude initialisation, controlling the spacecraft using observations of reference stars across the star mapper and data from the attitude and orbit control system. The on-ground attitude estimation process continued and the difference between on-ground and on-board estimates of the Tait-Bryan error angles and gyro drifts were monitored along with the on-board innovations to allow the on-ground support team to assess the convergence of the on-board algorithms. When the initial Tait-Bryan uplink was large (of the order of $0^\circ.5$) then it was necessary to wait until the on-board controller had recovered the initial attitude error and then fine-tune by uplinking a further set of Tait-Bryan error angles before switching control to the on-board algorithm.

Once real-time attitude determination had been seen to converge on-board, the image dissector tube should be capable of tracking stars. The image dissector tube was the primary detector for the scientific data, and its correct operation was an indication of real-time attitude determination convergence. Thus once the attitude and the gyro drifts were known on-ground, the values could be uplinked to the spacecraft for use by real-time attitude determination and normal operations could continue. In this way, the attitude was accurately known and the gyro drifts could be continually calibrated.

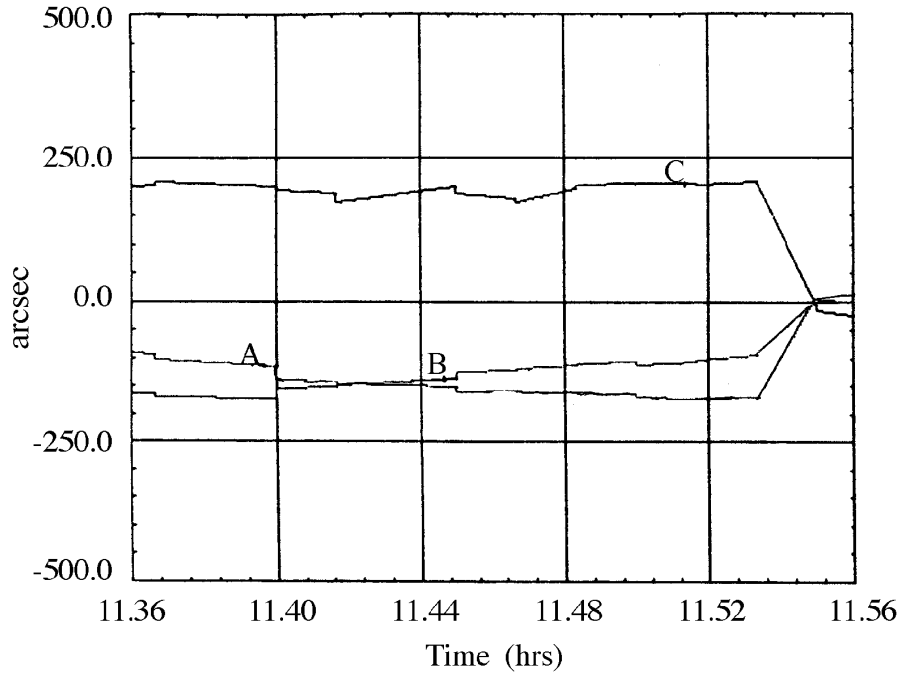

Figure 14.13. *Tait-Bryan angle uplink. The figure shows three lines denoting the differences between the on-board and on-ground estimates of the Tait-Bryan error angles ϕ, θ, ψ (lines A, B, C respectively). All three curves reduce to zero at uplink (to the right of the diagram).*

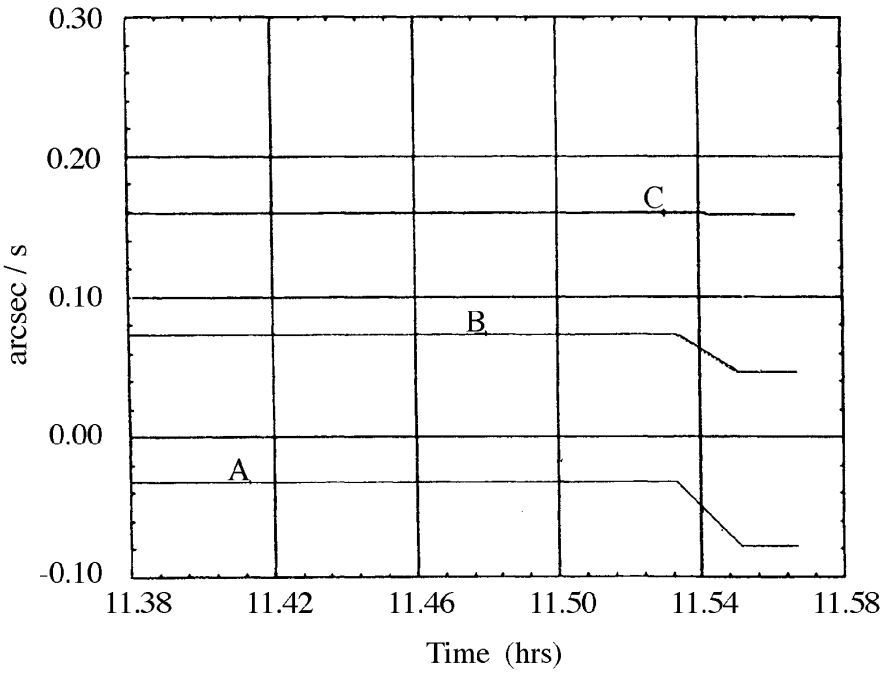

Figure 14.14. *Gyro drift relative to the spacecraft axes, as used on-board, for the x, y and z drifts (labelled A, B, C respectively). Values change after the uplink of new values from ground.*

14.4. Real-Time Attitude Determination Performance

In addition to ground real-time attitude determination, software ran continuously to monitor the attitude and the behaviour of the gyros and thrusters. Various alarms were to be triggered, on-board and on-ground, when an attitude anomaly occurred, such as a failed gyro or jammed thruster. If the spacecraft was visible from one of the ground stations, procedures could be implemented to minimise attitude drift. On-board, there were other safety measures, which could be triggered automatically or manually from ground, to protect the spacecraft should it start to drift for any reason. A sun sensor was used to protect the spacecraft from drifting more than 3° off the nominal scanning law with the immediate danger of the sun shining directly into the telescope, damaging the grid or detectors. If the spacecraft drifted past this threshold a so called emergency sun reacquisition order was triggered on board. The spacecraft then immediately slewed back to sun-pointing.

Data Recovery Percentages

The revised orbit induced real-time attitude determination divergence after approximately half of all perigee passes. The requirement to re-initialise real-time attitude determination prompted the decision to introduce a shift team of flight dynamics specialists to support spacecraft operations. Real-time attitude determination performance was continually monitored using the image dissector tube observations as the definitive criterion for convergence, i.e. if stars were being consistently observed by the image dissector tube, then real-time attitude determination must be converged. Once real-time attitude determination diverged (most often seen immediately after perigee) re-initialisation would take place using the ground real-time attitude determination software.

The overall percentage of science data achieved over the mission is shown in Figure 14.15. The figures are derived from ESOC's ground-based payload monitoring software (described in Chapter 10). 'Good' data refers to periods when stars were consistently observed by the image dissector tube during coverage periods. 'Bad' data refers to periods when real-time attitude determination diverged (usually immediately after perigee) during coverage periods. 'No data' refers to periods without ground station or computer support. The hibernation periods refer to those times when the spacecraft was spun up between three- and two-gyro phases and between two- and zero-gyro phases.

Once all three ground stations were made available, 'good' data recovery was around 71 per cent during the three-gyro operations (71 per cent of 24 hours per day). 'Good' data was recovered 87 per cent of the coverage time.

The time to perform real-time attitude determination re-initialisation varied greatly. This was due to two primary factors: (a) star mapper noise in the van Allen belts and (b) occultations after acquisition of signal.

The star mapper was particularly sensitive to background radiation in the van Allen belts. When highly active, after a solar flare, the star mapper could be blinded for

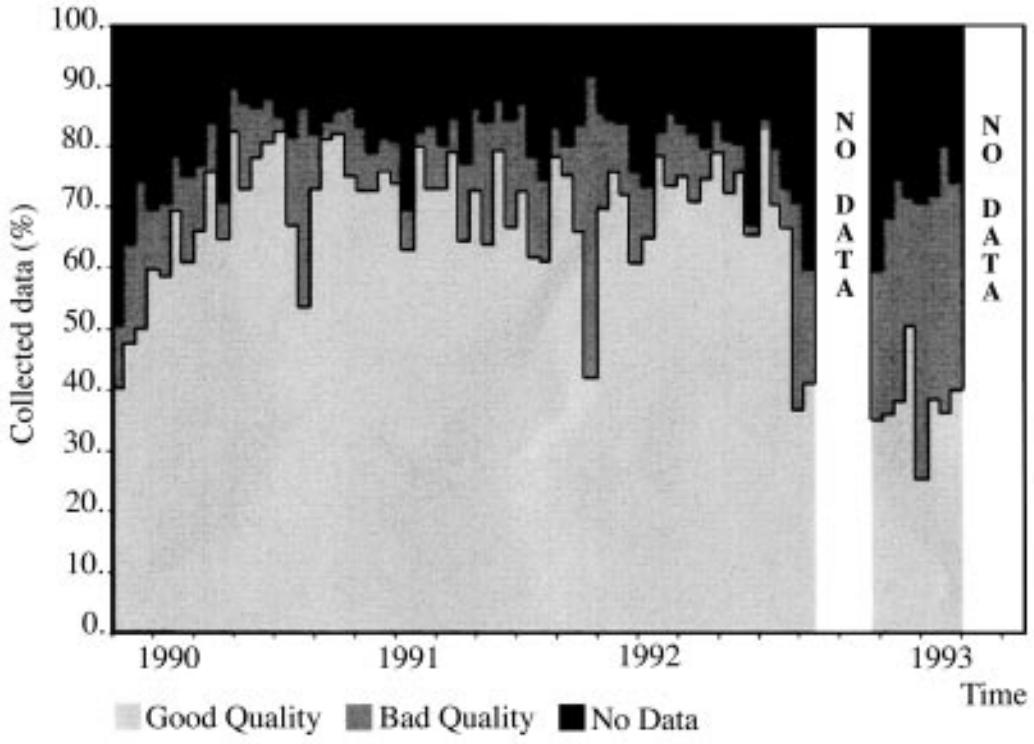

Figure 14.15. *Science data recovery as measured at ESOC.*

several hours per orbit due to the high background count rate. In the most extreme cases of solar activity, the star mapper would be blinded for the entire orbit, making convergence of real-time attitude determination impossible. Blinding of the star mapper also affected the ground real-time attitude determination software which filtered the star mapper data stream. Only once the noise was sufficiently low was it profitable to start real-time attitude determination re-initialisation procedures.

Seasons occurred where real-time attitude determination convergence could be recovered by switching to extended window after perigee. These seasons could last days or, latterly in spring 1992, weeks. At other times, real-time attitude determination would require initialisation every pass. These seasons were however unpredictable in their occurrence and were the subject of studies performed at ESOC and Matra Marconi Space. Resulting from these studies and input from the data reduction consortia, a significant number of changes (such as improved torque modelling and thruster calibration) were made to on-board and on-ground software which over the course of the mission improved the accuracy of real-time attitude determination in the apogee region and the chances of maintaining convergence though perigee. This had a direct consequence of improving the quantity and quality of the science data.

Typically at the start of 1990, the real-time attitude determination innovations (average difference between predicted and actual transit times converted to angular separations) were between 1 and 5 arcsec. By the beginning of 1992, the innovations were consistently around 1 arcsec, although the size began to increase again with the deterioration of gyro 5.

The principal improvements are described in the subsequent sections.

Central On-Board Software Patch and Parameter Tuning

As described in Chapter 9, in January 1990, the first calibrated grid rotation from the FAST Consortium was used on-board resulting in a significant improvement in on-board real-time attitude determination performance. The resulting innovations, which give a measure of the attitude corrections necessary between star mapper transits, were reduced from an average distribution of around ±5 arcsec, to around ±2 arcsec.

In March 1990, a new version (4.02) of the central on-board software was uplinked to the spacecraft. The new patch contained the following changes:

(a) extensions of the time-tagged buffer with an increase of allocated space from 640 words to 1504 words;

(b) extension of the memory area for the star data blocks which was increased from 1983 star data blocks (11898 words) to 2314 star data blocks (13884 words);

(c) improved star mapper filtering to eliminate false star mapper identifications in high noise.

The first change was necessary to perform efficient operations through long ground station outage. The second change allowed a denser selection of real-time attitude determination reference stars through long ground station outage (the 'sparse programme star file' described in Chapter 8). Clearly, the likelihood of maintaining real-time attitude determination convergence was proportional to the frequency of good star mapper measurements.

The third change was necessary since it became clear that high star mapper noise during van Allen belt passage was occasionally causing the on-board software to make a 'false star' detection. This erroneous measurement could be seen as real-time attitude determination as a larger than average innovation. Although the largest errors would be trapped by setting an appropriate innovation rejection threshold on-board, moderate sized errors were still apparent leading to increased noise in the overall real-time attitude determination performance. After the on-board patch was included, a clear reduction in the largest innovations was seen. Moreover, it became possible to make better use of the 'extended window' innovation threshold of 30 arcsec during high noise to help re-converge real-time attitude determination.

In May 1990, a swap from the B_T to the V_T channel for star mapper processing was made. This was due to the higher response in the V_T channel giving a higher number of successfully processed stars.

Over the following months, other tuning of on-board software parameters associated with star mapper rejection criteria was made. Statistics of reasons of star mapper measurement rejection were compiled over several weeks, correlating them to background noise and magnitudes of the reference stars. In this way, certain magnitude dependent parameters relating to expected height and width of the filtered signal were tuned. The result was a steady increase in the frequency of successful star mapper measurements for use by the real-time attitude determination system.

Reference Star Rejection Criteria

It was observed early on in the mission that the faintest real-time attitude determination reference stars were invariably rejected in the star mapper processing (95 per cent). Therefore, in March 1990, the limiting magnitude of such reference stars (flagged in the programme star file) was reduced from 10 to 9 mag. The result was a reduction in the load on central on-board software processing, allowing more of the brighter stars to be successfully processed.

Hipparcos Input Catalogue Updates

As described in Chapter 9, the NDAC Consortium was able to feed back their preliminary star positions and magnitudes to ESOC, via the INCA Consortium, to be used for the on-board programme star file.

In February 1991, Version 8 of the Hipparcos Input Catalogue was received and implemented at ESOC, containing approximately half the programme stars with positions updated from the preliminary data reductions. By November 1992, Version 10 was received, with nearly all of the programme stars containing updated results. The improvement in accuracy of the positions was, by that time, superior to the intrinsic performance accuracy of the real-time attitude determination.

This and the greater reliability in the magnitude estimates allowed ESOC to tune the selection criteria for attitude guide stars and to tune the on-board star mapper processing parameters to optimise the correct identification of guide stars in the star mapper data stream and virtually eliminate any possibility of mis-identifying stars, which could lead to an unexpectedly large innovation error in the on-board control.

Occultations and Eclipses

An occultation occurred when the Earth or Moon crossed one of the fields of view of the telescope (the definition of the nominal scanning law prevented the Sun from ever doing so). Light from either body could have damaged the sensitive detectors, hence shutters were closed to protect them. Unfortunately that meant that no stars could be seen by either the image dissector tube or star mapper and real-time attitude knowledge could deteriorate, resulting ultimately in divergence.

Occultations by the Earth or the Moon required that the star mapper shutters be closed for up to 30 min at a time, near the perigee region. These would often hinder ground real-time attitude determination re-initialisation procedures. At times, it would be necessary to restart the procedure, if the occultation interrupted the star pattern recognition before convergence of the ground software could be achieved.

In 1991, the mission planning software at ESOC was modified to allow the shutters to remain open during occultations whenever a dark area of the Earth's surface occulted one or both fields of view. Since there was no risk to the instruments from the dark side of the Earth, real-time attitude determination was more likely to be maintained using star measurements from the other field of view. This was particularly valuable at

lower altitudes, where the Earth's disc was much larger and the occultation duration consequently longer.

Gyro Projection Matrices

As described in Chapter 9, in July 1991, the NDAC Consortium was able to supply a newly calibrated gyro projection matrix for uplink to the spacecraft, to account for gyro misalignments with respect to the nominal positions.

Gyro Drift Monitoring and Correction

As described in Chapter 13, from March 1992 onwards, ground procedures were changed to allow the resetting of the on-board drift estimates as soon as possible after perigee to the last values observed before perigee. Subsequent careful monitoring for drift evolution was then performed throughout the orbit to ensure that the best drift estimates were maintained.