4. SATELLITE DATA PROCESSING

The Tycho data processing chain started with two very large processes: the 'prediction of group crossings', and the 'detection of transits', treating the raw data from the satellite, as delivered by the European Space Operations Centre (ESOC). The general role of the two processes was outlined in Chapter 2. They are described in more detail in this chapter.

4.1. Prediction of Group Crossings

The prediction process was executed at Astronomisches Rechen-Institut, Heidelberg. Its task was to calculate expected times of group crossings of celestial objects from input data giving:

- the attitude (i.e. pointing on the sky) of the satellite as a function of time;
- the orbital position and velocity vectors of the satellite as functions of time;
- the apparent positions of the celestial objects, depending on time and on the orbital position and velocity of the satellite;
- the precise optical geometry of the telescope and grid assembly, i.e. the geometric instrument calibration.

A 'group' is a set of four parallel star mapper slits in the focal plane of the Hipparcos telescope (see Figure 1.1). A group crossing is the transit of a star's image across a fiducial line defining the centre of gravity of the four slits in the direction of the motion of the star images across the focal plane. This fiducial line is called the centre line of the group.

The celestial objects to be treated by the prediction process included the 3.15 million entries (mainly stars) of the Tycho Input Catalogue, plus 59 solar system objects: Venus, Mars, Jupiter, Saturn, Uranus, Neptune, the 4 Galilean moons, Titan, and the 48 brightest Minor Planets.

The actual Tycho data processing chain contained several iterative steps of 'prediction' and 'prediction updating'. The present section is restricted to what in Chapter 2 was called 'first prediction'. The various steps of iteration will be described in Chapters 5, 6 and 9.

Basic Algorithm

The basic task of the prediction process is simple. The attitude of the satellite is a time-dependent rotation matrix, transforming celestial coordinates into field coordinates in the Hipparcos focal plane. It is given in some tabular form, with the entries of the table being valid for a dense set of (spacecraft-defined) instants of time. Two such instants (t_1, t_2) define a short time interval (1.067 s in the case of 'first prediction') within which the attitude can be linearly interpolated with sufficient accuracy. For any such time interval, prediction does the following:

- get attitude and satellite status data for t_1 and t_2. Use the satellite status data to decide whether scientifically useful observations may possibly be performed during the interval (t_1, t_2). If not, skip the interval;
- select all celestial objects that may experience a group crossing during the interval (t_1, t_2) in either of the two fields of view of the telescope (this is the most complex part of the software);
- calculate the satellitocentric apparent positions of these objects for $(t_1 + t_2)/2$, i.e. the interval mid-time (this is the most time-consuming part of the software);
- transform the apparent positions into field coordinates in the focal plane of the Hipparcos telescope, for times t_1 and t_2, using the corresponding attitude matrices;
- check (for each object) whether it has crossed the centre line of any of the slit groups between t_1 and t_2. If so, compute the predicted transit time by linear interpolation of the field coordinates;
- for each transit found, compute auxiliary data (such as the instantaneous scan speed, the instantaneous direction of scanning, the field coordinates at transit time, jet firings of the attitude control system etc.) and create the output records for the PGC ('predicted group crossing') data stream. There is one such record per predicted group crossing;
- sort the records according to the predicted transit time;
- write everything to magnetic tape.

In spite of this basically simple task, the actual prediction software was large and complex. Its complexity was enforced by the sheer computational size of the task. The selection of the objects to be treated in each time interval (see second item above) was a complicated multi-stage process. It had to be very efficient, since it decided on the number of apparent positions to be computed. The computation of one high-precision apparent position for a star took about 1 ms on the IBM mainframe computer at Astronomisches Rechen-Institut, Heidelberg.

Thus, even if only 10^{-4} of all Tycho Input Catalogue entries would have been considered in each time interval, the prediction processing would have been slower than the satellite observations. Likewise, the actual retrieval of the selected objects from the magnetic-disk Tycho Input Catalogue file had to be organized very efficiently. This was necessary because, from the point of view of catalogue access, Hipparcos scanned the sky in a highly erratic fashion.

Input Data

Prediction had to treat several types of input data, which will be described very briefly in this section. The largest amount of input data was provided by the European Space Operations Centre (ESOC), in the form of roughly 1300 magnetic tapes with raw telemetry from Hipparcos. Slightly more than one big tape reel (2400 feet), corresponding to about 150 Megabytes of data, was produced by the satellite on each day of scientific operations. The sum of all other types of input data together amounted to less than one per cent of this volume.

The real-time satellite attitude was provided in the so-called 'attitude and orbit control system' files on the telemetry tapes. The orientation of Hipparcos in space was given in the form of three 'error angles' relative to the nominal scanning law. The nominal scanning law itself was not tabulated on the tapes, since its definition was part of the mission specification. The rotation matrices mentioned in the previous section were derived from the attitude and orbit data.

A variety of satellite status data were also derived from the attitude and orbit data files and from the so-called Data Catalogue file on each telemetry tape. The relevant data included gas jet firings, shutter and electronic switch positions, orbital correction manoeuvres, bad telemetry signal levels, the two free parameters of the nominal scanning law, etc. Some of these were used to decide on the usefulness of the telemetry (i.e. to switch the prediction process on and off), and some of them were just copied to the output stream as warnings to the users.

Geocentric orbital ephemerides of the Hipparcos satellite were provided by ESOC, as a separate file on each of the telemetry tapes. Instantaneous position and velocity vectors of Hipparcos were computed from a Chebychev polynomial representation of its orbital motion.

Different sorts of ephemerides had to be used for different groups of solar system objects. The standard Development Ephemeris DE200 of the Jet Propulsion Laboratory (Pasadena, USA) was used for the major planets, including Earth. Special ephemerides on *ad hoc* formats were provided for the minor planets and for the moons of the major planets by A. Bec-Borsenberger and J. Arlot, respectively, of Bureau des Longitudes, Paris. More details on these and on their usage are given in Section 10.4.

The primary inputs provided by the Tycho Input Catalogue were the mean positions, proper motions and parallaxes for the computation of apparent positions of stars. Besides these, the Tycho Input Catalogue provided magnitudes, flags etc., which were simply copied to the output data stream of the prediction process. In addition to the 3.15 million Tycho Input Catalogue objects, a list of almost 100 000 so-called serendipity points was used for the serendipity process, as described in Section 5.5.

Last but not least, prediction had to make use of a precise model of the optical geometry of the telescope and grid assembly. For Tycho, this model essentially consisted of the field coordinates of the centre lines of the various slit groups in the Hipparcos focal plane, plus the precise value of the basic angle between the two fields of view. The pre-flight calibration provided by ESOC turned out to be too imprecise for prediction. Thus, early in-orbit calibrations produced by the FAST consortium were used throughout the mission.

The specifications of all input data except the Tycho Input Catalogue and DE200 are unpublished, as are the data themselves. The telemetry data were defined in the ESOC 'Data Delivery Interface Document for Hipparcos'. The interpretation of the satellite orbit file is described in an additional letter by ESOC's Attitude and Orbit Division. The ephemerides for minor planets and moons are described in letters from their originators only. The FAST instrument geometry description is defined in an internal document called 'FAST Calibration Document', and that of NDAC in NDAC-internal interface documents titled 'Great-Circle Interface Document 7' and 'C-238'. These facts are mentioned here in order to explain the absence of references to published material in the present section.

Output Data

The output of 'first prediction' consisted mainly of the Predicted Group Crossing (PGC) data stream, indicated in Figure 1.2, plus some protocol and log files. The PGC data stream was used in the 'detection' process and the various 'prediction updating' processes. It mainly contained time-ordered PGC records, one for each predicted group crossing. In addition, there was a header record every 10.66... s of operational mission and a Data Catalogue file roughly every day of mission, both giving auxiliary data on satellite status, timing, orbit etc. The precise description of the data stream is given in a Tycho-internal booklet under the title 'Tycho Interface Document'. The total data volume of the PGC data stream for the entire mission was roughly 200 magnetic tapes of 2400 feet each, corresponding to about 30 Gigabytes.

Computation of Apparent Positions

Satellitocentric apparent positions of celestial objects were computed to an accuracy of better than 1 mas. The relevant physical effects included in the calculations are, for stars: proper motion, stellar aberration (due to the orbital motion of the Earth around the Sun, and of the satellite around the Earth), annual parallax, relativistic light bending by the Sun and the Earth (but not the other planets; they are irrelevant for Tycho). The astronomical and mathematical basis of the procedure was described by Walter *et al.* (1986).

Prediction used the following model for proper motions: The mean positions given in the Tycho Input Catalogue (TIC) were used in the form of cartesian position vectors (x, y, z). These referred to a fixed epoch t_{TIC} = J1990.0. The Input Catalogue proper motion components (in right ascension and declination) were used in the form of time derivatives (x', y', z') of the cartesian position vector. The position vector of an object at a given epoch t was calculated in two steps. First an intermediate vector $(x, y, z) + (x', y', z') \times (t - t_{\mathrm{TIC}})$ was computed. Then this vector was normalized to unit length to give the position vector of the object at epoch t. Consistent application of this proper motion model was necessary in the 'astrometry' task, especially for nearby stars with large proper motions and for stars close to the celestial poles.

Retrieval of Tycho Input Catalogue Objects from Magnetic Disk

Within an attitude time interval (t_1, t_2), the centre lines of the slit groups in each field of view swept across a part of the celestial sphere having the shape of a symmetric right-angled pentagon (see Figure 4.1). Only objects lying within that part of the sky

could possibly produce a group crossing within the given time interval. For each time interval there were exactly two such pentagons on the sky: one for each field of view. Their location and orientation on the sphere were determined by the satellite attitude at time t_1; their length in the direction of scan was determined by $t_2 - t_1$. Each of the pentagons comfortably fits into a circle of radius $r = 23$ arcmin (for $t_2 - t_1 = 1.067$ or 2.13 s) around a central point P on the line of symmetry (see Figure 4.1). This fact was utilized for efficient TIC access in the following way.

The celestial sphere was divided into small 'regions' in a special way. A running number n was assigned to each of these regions. For any point on the sky (e.g. a Tycho Input Catalogue object or the central point P), this region number can be determined by a simple algorithm. The main point is the following (see Figure 4.2): For any position on the sky, situated in an arbitrary region n, the circle of a fixed radius s around that position is entirely contained in that region n and the immediately adjacent regions (shaded in Figure 4.2). The numbers of the adjacent regions can again be determined by a simple algorithm. The 'characteristic size' s of the regions is chosen slightly larger than the above-defined radius r of the pentagon.

To utilize this concept the Tycho Input Catalogue file on magnetic disk was organized and indexed according to these regions on the sky. Retrieval of the Tycho Input Catalogue objects necessary for one field of view in a time interval (t_1, t_2) started with the computation of the region number n for the central point P of the pentagon, and the region numbers of all the immediately adjacent regions. Then all objects within these regions (typically a few hundred to a few thousand) were read into computer main storage. But these objects were still too many for the computation of apparent positions. The next selection step, therefore, simply transformed the Tycho Input Catalogue positions to field coordinates (using the attitude matrix for time t_1), and kept only those objects situated close to the pentagon or its forward extension (bold dashed lines in Figure 4.2).

Once this retrieval procedure had been executed for the time interval (t_1, t_2) it did not need to be repeated as long as the central point P remained situated in the same region during subsequent time intervals. For a reasonable 'characteristic size' s this condition held, on average, for about 10 consecutive time intervals of 1.067 s each. During this time, the selection of objects was permanently refined in order to save even more computation of apparent positions. For instance, every object having crossed the vertical slit group was removed from the selection for following time intervals. It could not produce another group crossing during this particular transit across the field of view.

The sole purpose of this mechanism, which occupies a large fraction of the entire prediction program code, was the saving of computing time for apparent positions. Just for the sake of completeness this paragraph is concluded with the recipe for the division of the sphere into regions having a 'characteristic size' s:

- let the regions be rectangles in (α, δ) space. Choose their size in the direction of δ (i.e. declination) as $s = \pi/(2N)$, where N is an integer, such that s is slightly larger than r;
- take as region number 1 a circle around one pole with radius s (this is also a rectangle in (α, δ) space);
- add circular declination zones, each having width s;
- divide the first such zone into six regions of identical shape, with one of the north-south boundaries being at $\alpha = 0$;

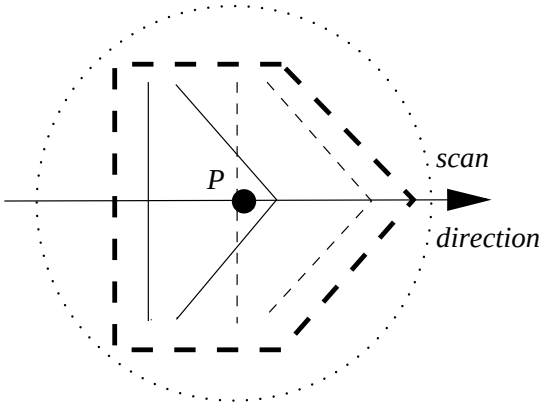

Figure 4.1. *The scanning pentagon on the sky. The figure shows the centre lines of the various slit groups projected onto the celestial sphere, for two instants of time (in one of the two fields of view only). Their location at time t_1 is given as thin full lines, that at time t_2 as thin broken lines. The horizontal shift between the full and the broken lines is equal to the time difference $(t_2 - t_1)$, multiplied by the scanning velocity v of the satellite. The bold dashed lines indicate the symmetric right-angled pentagon which is discussed in the text. Its central point P is indicated by the filled circle. The pentagon contains all celestial objects that can possibly produce a group crossing in the time interval (t_1, t_2) for the particular field of view. The enclosing circle, which is discussed in the text, is also shown (dotted).*

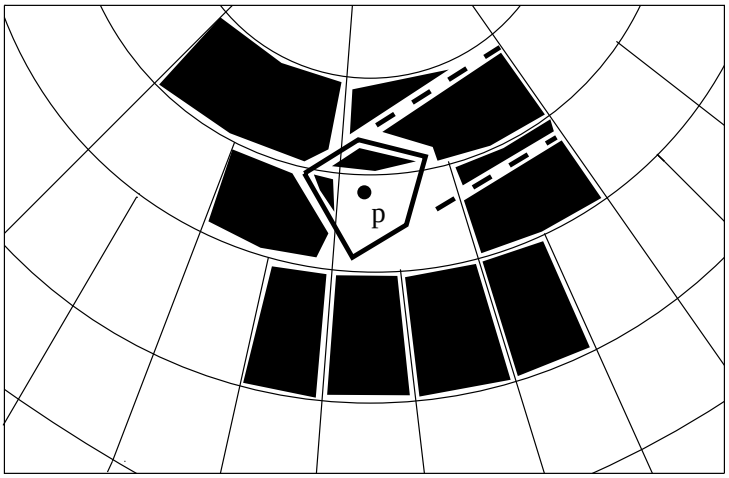

Figure 4.2. *Principle of Tycho Input Catalogue access. The figure shows a part of the celestial sphere, divided into the 'regions' (thin lines), as described in the text. The pentagon for a given instant of time is shown in bold lines, the corresponding central point P as a filled circle. The region containing P is evident. 'Immediately adjacent' regions are those having a common corner or boundary line with this first one. They are shaded in the figure. The 'forward extension' of the pentagon is outlined by dashed lines.*

- divide all the remaining zones analogously into 6×2^m regions, where m is an integer chosen such that:
  - the regions are as small as possible, but
  - the shortest distance between any two rectangle 'corners' is at least s.
- assign a serial number to each individual region by starting at one pole with number 1, proceeding from zone to zone and numbering the regions within each zone according to the value of α at their centres.

This recipe leads to a simple algorithm for the computation of a region number from a given celestial position, and of the numbers of the immediately adjacent regions.

Solar System Objects

The treatment of the solar system objects in the prediction process was quite different from that of the Tycho Input Catalogue objects. Its description is deferred to the chapter on Tycho reprocessing (Section 10.4), since only reprocessing data was used for the solar system objects. The astronomical results on the solar system objects are described in Chapter 15.

Consistency between Geometric Instrument Calibration and Attitude

Prediction combined the real-time satellite attitude and a precise model of the instrument geometry to compute star transits across the star mapper slit groups. The attitude, in turn, was determined on board the satellite from actually observed transits—implicitly using some instrument geometry as input. The on-board instrument model was chosen by ESOC to optimize the satellite operations for the main grid. The Tycho prediction model, on the other hand, was chosen to give the best possible predictions for the star mappers. Any inconsistency between the two instrument models would necessarily lead to systematic offsets of the predicted crossing times from the actual crossings of stars. Such offsets were indeed observed during the mission.

Prediction was started using the best available pre-launch instrument geometry. The predicted group crossings were compared with actual crossings (as observed by Tycho's detection process). Very large systematic deviations were found, caused mainly by the big rotation of the grid assembly relative to its nominal position inside the telescope. That rotation was taken account of by an early in-orbit geometry ('Version 2') produced by the FAST consortium. A constant offset 'observed minus predicted' of about 2 arcsec remained, which was removed by an *ad hoc* correction to the FAST geometry.

Beginning with 23 January 1990 (i.e. 2 months after the start of scientific observations) ESOC implemented the grid rotation into the on-board instrument model. This greatly improved the precision of the real-time attitude, but again introduced large deviations in the predicted group crossings. A small part of these deviations was compensated by a change to an improved FAST geometry ('Version 3'). However, the major part was due to a systematic attitude offset of the kind discussed in the first paragraph of this subsection. Its character was such that it could not reasonably be remedied by another correction to the instrument geometry. Instead, it was removed by an *ad hoc* correction to the attitude. This gave satisfactory results throughout the mission. The remaining systematic deviations were smaller than 0.1 arcsec, which is irrelevant for the 'first prediction'. The *ad hoc* attitude correction consisted of a constant rotation

by −5.9 arcsec and −0.65 arcsec, respectively, around the satellite's x and z axes (for a definition of the axes see Volume 2, Section 8.1).

Actual Processing

Trial and verification runs started on 9 November 1989, when the first provisional telemetry tape with real Hipparcos data was delivered by ESOC. The final mass processing started in March 1991, and was completed in April 1994.

All computing was done at Astronomisches Rechen-Institut, Heidelberg, using the IBM-3090 computer of the computing centre of the University of Heidelberg.

Prediction Redoing

For most of the Hipparcos mission the real-time attitude provided by the satellite had its nominal precision, i.e. about 1 arcsec, which was sufficient for the purpose of 'first prediction'. This was not true, however, for the very first and very last parts of the mission. The attitude was much worse for data collected before 23 January 1990 (i.e. for 63 days of scientific operations) and after 7 July 1992 (169 days of scientific operations), with errors frequently exceeding 5 arcsec. The start of the mission was impaired by the initially unsatisfactory instrument geometry model and by still sub-optimal attitude control. The end of the mission suffered from the rapid deterioration of the z-axis gyro and, after its total failure, from the reduced attitude information provided by the two-gyro operational mode. However, during both time intervals Hipparcos did actually collect high-quality star mapper observations.

Bad real-time attitude meant the loss of part of these observations, simply because the 'predicted group crossing interval' (see the paragraph on detection and estimation in Section 2.4) extended only 6 arcsec on either side of the predicted group crossing epoch. Therefore it was decided that the bad parts of real-time attitude should be replaced by on-ground attitude for the purpose of first prediction. This decision created the task of 'prediction redoing', and analogous tasks for detection of transits, prediction updating and transit identification. The repeat of about 20 per cent of these processes was considered affordable within the agreed Tycho time schedule. Software capable of doing prediction with on-ground attitude had been planned in any case (for the 'reprocessing' step, see Chapter 10). Thus, the prediction redoing task was executed immediately after the reprocessing software and the necessary on-ground attitude became available in autumn 1994.

The usage of on-ground attitude instead of on-board real-time attitude made prediction redoing more precise and reliable than the original prediction, and in addition it made it also quicker, for reasons explained in Chapter 10. The on-ground attitude files did not provide some of the information contained in the telemetry tapes. This is why the data interfaces between the various Tycho processing steps had to be slightly modified for redoing. In consequence, the software for detection of transits and for prediction updating also had to be modified for the redoing step.

4.2. Background Determination

The detection and estimation process, including the background determination, is the very core of the Tycho data reduction scheme. Therefore it was described at some length in the overview in Chapter 2. This section and the following two sections give additional details.

Introduction

The raw Tycho signal is the sum of photon counts produced by transitting stars and a background. The quality and efficiency of the detection and estimation processes, and thus of the whole data reduction chain, depend strongly on the quality of the background determination. This is mainly because the signals of most of the star transits used by Tycho were not far above the background level.

Background determination was carried out during the initial data preparation for the detection process. It was the first piece of TDAC software looking into details of the Tycho raw data. In general, the method of background determination was independent of star transits and the normal data reduction. The background was derived and provided for any given stretch of data and for both channels of the Tycho experiment. Moreover, all background data were stored in a self-contained data base giving access to information about the satellite environment, the sky background and the zodiacal light during the mission for all data used by TDAC. The whole background data base together with some orbital data of the satellite is publicly available. It can be used for studies of the variations of the van Allen belts or of the influence of solar flares on the environment of satellites (see Wicenec & Bässgen 1992; Wicenec & van Leeuwen 1995).

Mathematical Principle

The background used by Tycho was based on the median value of a set of count rates, where the number of count rates in the set depended on the non-stellar disturbances in the raw counts (see below). If a set of count rates is sorted, the median is defined as the middle value in the case of an odd number of count rates. In the case of an even number of count rates the median is determined by interpolation between the two middle values (Spiegel 1961).

The very low detection limits used by the detection process (signal-to-noise ratio = 1.8 for transits outside the predicted group crossing intervals, and signal-to-noise ratio = 1.5 for transits inside these intervals) yield some problems due to disturbances by spikes (see Chapter 2) or dense star fields. These were circumvented by the usage of median (rather than mean) values for the definition of the background. Any set of count rates is a statistical representation of the distribution function F_d of the counts. Normally one would expect F_d to be a Poissonian distribution for the type of independent processes detected with the Tycho photomultipliers. However F_d is disturbed by stars and spikes. In the actual satellite data, both the background level and the strength of the deviations from Poissonian statistics varied strongly (see Figure 2.3).

Even the median value is sensitive to a long tail of the distribution function (Press *et al.* 1986), i.e. to an excessive asymmetry of F_d due to a large number of spikes or a high density of faint stars. Such effects slightly shift the derived background to higher values. This bias was deliberately accepted by Tycho. It was utilized in the detection process to limit the number of false transits in heavily disturbed formats.

Several statistical estimators in addition to the median were used to judge the disturbance of the distribution function and to decide how long the set of counts could be chosen for the background determination. These include the lower and upper quartile of the distribution function of the count rates in the set and the mean of the four maximum count rates. For a set size of 6400 samples, i.e. 10.66... s of data, the difference of the median background to one estimated from the lower quartile was less than 1 per cent (indicating very little disturbance) for about 90 per cent of the data. It increased to the order of 10 per cent for the highest star densities, where both estimates of the background had an internal uncertainty of more than 50 per cent. Faint stars of the Milky Way band which are not resolved by Tycho as well as all clusters of resolved stars caused a bias of the background level by disturbing the distribution function.

Format Background

The so-called 'format background' was derived in the manner described above from sets of 6400 samples each, corresponding to one telemetry format, or 10.66... s of data. Usually the background varied very slowly, being essentially constant over periods much longer than one telemetry format. Thus, for at least 75 per cent of the sky (low to moderate star densities) and 50 per cent of the orbit (apogee region), the format background was applicable.

Local Background

Whenever the condition of essentially constant background was violated, the so-called 'local background' was used instead of the format background. It was computed by exactly the same method as the latter, but with the size of the data set reduced to 192 samples, corresponding to 0.32 s of time, or 54 arcsec on the sky. Cases where the local background had to be used included:

- transits of very bright stars ($m < 5$), of dense open star clusters and of dense Milky Way regions;
- scanning of the zodiacal light;
- crossing of the van Allen belts and strong variations of the energetic-particle background due to magnetic storms, solar flares etc.;
- straylight produced in the instrument by a very bright extended object.

The variation time scales in these cases ranged from some 100 samples to several formats. Careful studies of Tycho data showed that the shortest background variation time scales were due to bright stars and internal reflections where the latter ones had their source in very bright extended objects such as Venus, Jupiter or Saturn and therefore occurred very seldom. Figure 4.3 shows the lower part of a transit of a star with $B_T = V_T = 2.2$ mag after the nonlinear folding of the detection process (see Section 4.3). The very wide extension of the 'wings' of the main peak may disturb transits with a separation up to approximately ±150 arcsec, or about 500 samples, in scan direction. If

the format background would have been used in these cases the detection process would have produced lots of 'false' and too bright transits respectively, because the amplitude estimation for transits on such wings would result in an overestimation of the signal due to the local excess of the background. This effect was detected for stars down to $B_T \simeq 5$ mag and for open clusters. The local background was used therefore, in order to avoid 'holes' around bright stars in the Tycho Catalogue.

The length of 192 samples for the definition of the local background was chosen because about 200 samples cover the main peak of a bright star. The usage of the local background yielded accurate results for the cases of bright stars, high star densities and bright extended objects. Figure 4.4 shows the count rates of one telemetry format (filling the white area) and the local background derived for each sample of that format (black line). The variation of the local background exceeds ±10 counts, while the format background is about 16 counts. The variations are due to many (mostly faint) star transits while both fields of view are scanning the galactic plane. The number of predicted transits is around 600 in this format.

Tycho Background Features

The van Allen belts produced the most prominent features in the background data (see Figure 4.5). They appear as high steep peaks above the 'level' of the apogee background. Normally there should be four peaks during every orbit, but the two inner peaks are not complete, because of the telemetry loss during perigee. In the apogee region one may recognize numerous little peaks. They are due to the crossings of one or both of the telescopes's fields of view of the galactic plane, other regions of high star density or the zodiacal light.

The apogee zone, which was most interesting for Tycho, was not free of disturbance. The combined background level in the B_T and V_T channels, (called $b_{B_T+V_T}$ in the following) varied from 5 to 30 counts per sample. The time interval around the apogee which was not affected by the van Allen belt slopes, varied between some 10 minutes and about 5 hours. This duration was directly correlated with the height above the Earth where the van Allen belts begin to disturb the background. The maximum height was nearly equal to the apogee height of the satellite (36 000 km above ground), whereas the minimum height lay around 20 000 km. The variations within the undisturbed apogee zone amounted to approximately ±3 counts per sample with a mean value of about 6 counts per sample in $b_{B_T+V_T}$.

The peaks of the outer van Allen belts sometimes showed a 'bifurcation' (see Figure 4.6). This effect has been reported by Frank *et al.* (1964). It is due to a series of geomagnetic storms where the effects of a later one is superimposed on the previous, resulting in a doubling of the outer van Allen belt. These effects were known to last for a couple of days, which is in agreement with the Tycho data.

Figure 4.7 shows a different representation of the $b_{B_T+V_T}$ data. The count rates are given in counts per second (logarithmic scaling) versus the length of the satellite's location vector in units of the Earth radius (R_e). The different locations of the minima between 2 and 3 R_e indicate the passage of the satellite at the day and night side, respectively. The minima between the two belts sometimes were nearly as low as the count rates in the apogee region. However, studies of the spike density showed that the data between the belts had to be treated very carefully, and data in the range of the van Allen belts were not usable for TDAC at all. Moreover it turned out that the height of the satellite

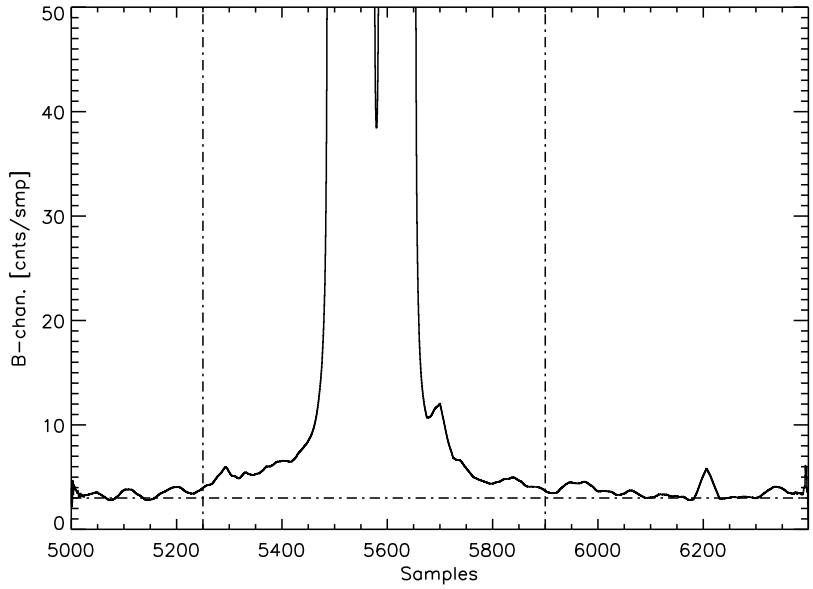

Figure 4.3. *Lower part of the count rates of a bright star transit. The wide wings of the instrument response function extend to about ±500 samples, corresponding to about ±150 arcsec on the sky.*

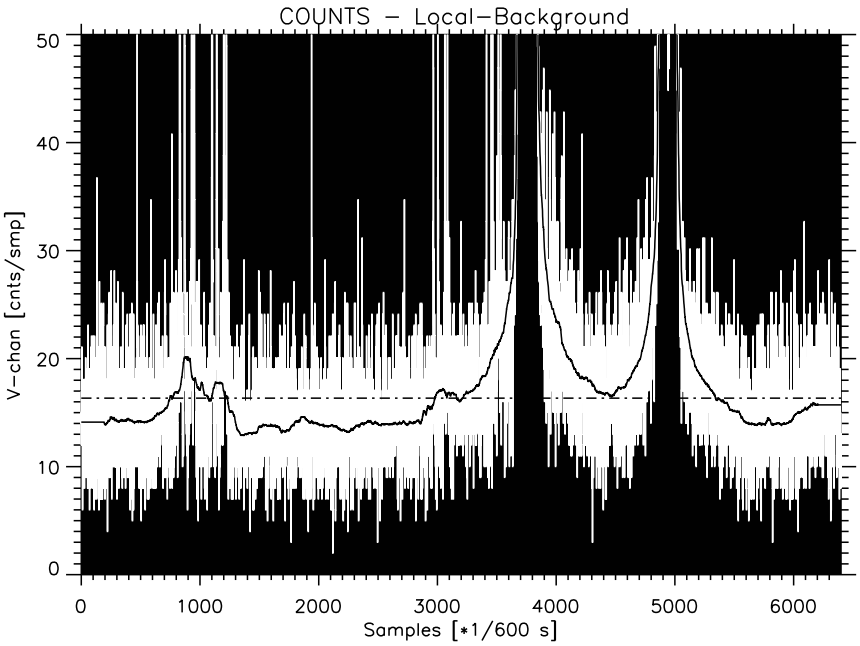

Figure 4.4. *Count rates and local background of one format of Tycho data where both fields of view are scanning in the galactic plane. Full line: local background; dash-dotted line: format background; white area: individual count rates.*

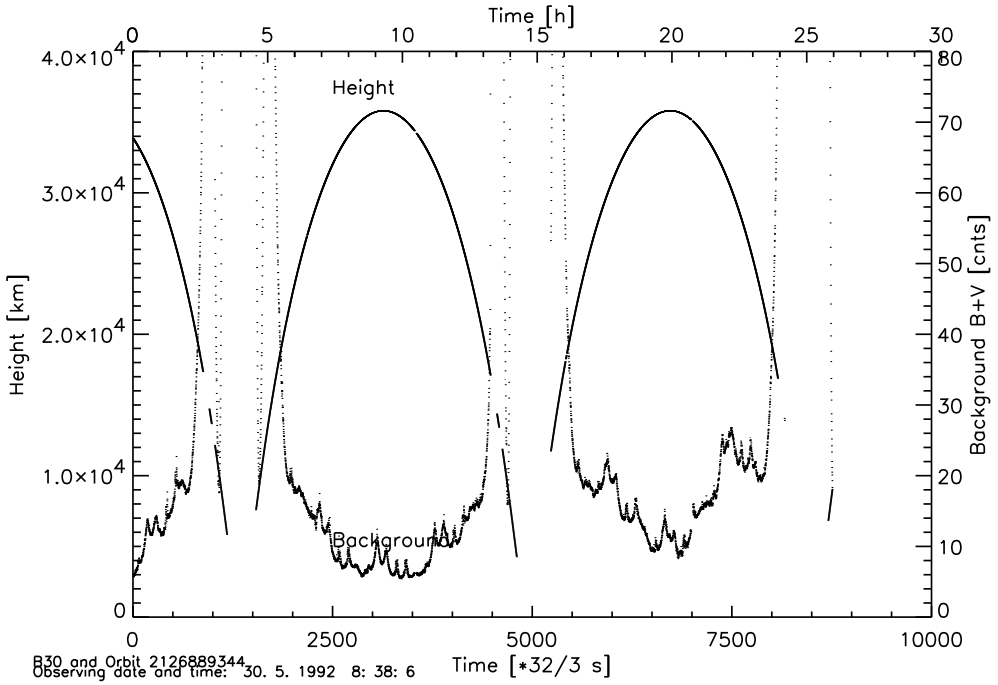

Figure 4.5. *The Tycho background $b_{B_T+V_T}$ level (small dots) compared to the height of the satellite above ground, for about 2 successive orbits. The individual features are discussed in the text.*

alone could not be used as an indicator of 'good' data. This was made impossible by the strong variations of the background level at apogee and of the height of the disturbing belts due to the solar activity.

4.3. Detection of Star Transits

The task of detection was to find signals due to any star-like object on the sky in the Tycho data stream, and to estimate the precise observed group crossing times and amplitudes of the signals from input data giving:

- the predicted group crossing times and auxiliary data from prediction;
- the raw photon counts from the B_T and the V_T channel of the Tycho experiment, called the Tycho data stream hereafter;
- calibration data.

Most of the complexity of the Tycho data reduction scheme stemmed from the very low limits used to detect predicted signals from starlike objects. The limits were chosen and adjusted in a way to allow about one or two false detections per one true detection. A detection is fully described by a transit time, the signal amplitudes, the backgrounds and the signal-to-noise ratios in both channels, accompanied by some auxiliary data which originated in the Tycho Input Catalogue (TIC), in the prediction process and in the Tycho data stream.

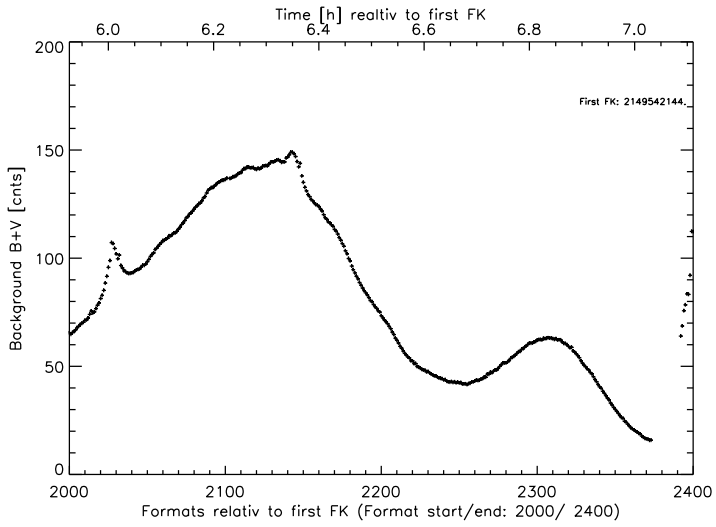

Figure 4.6. *One peak of the outer van Allen belt showing the bifurcation effect referred to in the text. The small peaks at the extreme left and on top of the main peak are caused by crossings of the galactic plane by the preceding and following field of view respectively. The scale along the top of the figure gives the time in hours with respect to an arbitrary origin, the bottom scale is an 'internal' timescale.*

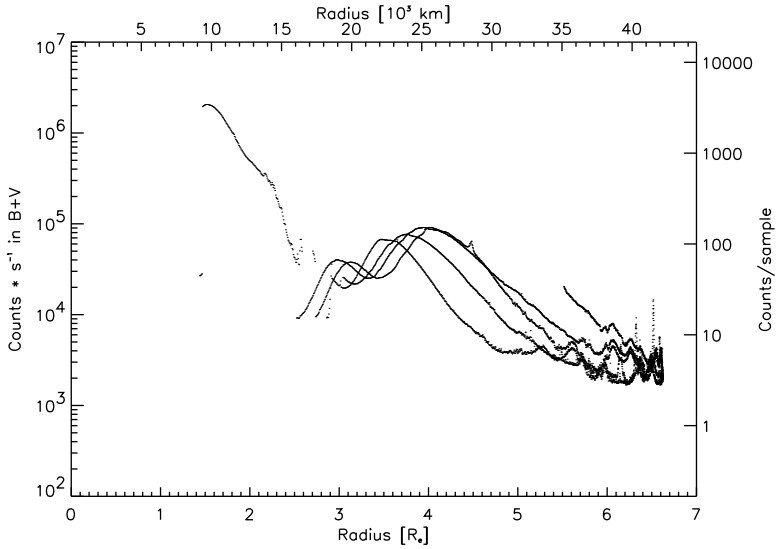

Figure 4.7. *The background count rates as a function of the satellite's distance from the centre of the Earth, for several consecutive orbits. The individual features are discussed in the text.*

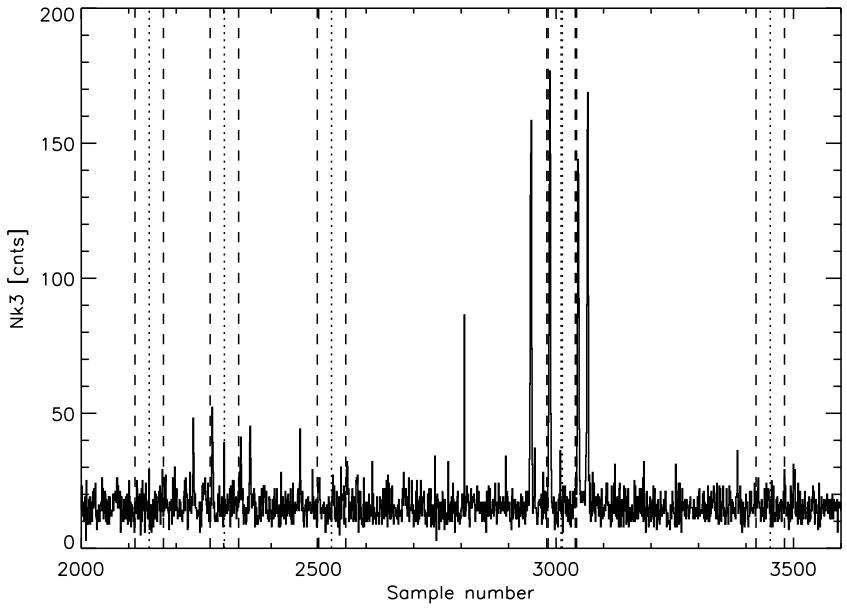

Figure 4.8. *Raw photon count rates (N_{k3}) in B_T + V_T. The stretch of data covers about 450 arcsec in scan direction, where some brighter stars are visible. For stars brighter than 10 mag in either of the two colours, the Predicted Group Crossing (dotted) and the corresponding Predicted Group Crossing intervals (dashed) are shown.*

Input Data Bases

The Tycho data stream contained the photon count rates in B_T and V_T, sampled at 600 Hz, plus information about the timing of the observations, and flags describing their quality. The photon count rates were given in compressed one-byte values in a semi-logarithmic form. They were decompressed after reading.

The Predicted Group Crossing (PGC) data stream gave, for each predicted group crossing: the identification of the Tycho Input Catalogue entry concerned; the PGC epoch; the type of the slit group (inclined or vertical); the field of view (preceding or following) in which the object was expected to appear; the vertical coordinate on the slits, and the actual scanning velocities (or rather the image velocities of the star) along the nominal scanning direction and its orthogonal direction, plus information originating from the Tycho Input Catalogue, e.g. ground-based magnitudes in B and V. Figure 4.8 shows an extracted part of raw photon counts with the location of some predicted group crossings for some transits, and the corresponding predicted group crossing intervals.

Calibration data included the decompression tables for the photon count rates, the geometric calibration of the star mapper slit groups, and single-slit response functions for the different slit groups (vertical/inclined, upper/lower half of slit, preceding/following field of view, B_T/V_T channel). The detection process used the photon count rate decompression tables given by ESOC, and the geometric star mapper calibration and single-slit response functions derived by the NDAC team at the Royal Greenwich Observatory.

Output Requirements

The subsequent tasks: recognition (Chapter 5), astrometry (Chapter 7) and photometry (Chapters 8 and 9), required the following from detection and estimation:

- background values in B_T and V_T;
- transit time estimates in B_T, V_T and $B_T + V_T$;
- B_T, V_T and $B_T + V_T$ signal amplitude estimates;
- signal-to-noise ratio for each detected transit;
- auxiliary data copied from the Predicted Group Crossing and Tycho signal data streams and from the Tycho Input Catalogue.

Separate B_T, V_T and $B_T + V_T$ transit time and amplitude estimates were obtained using the B_T and V_T count rates and the sum of both. The times were the result of an inverse interpolation in the folded counts, the amplitudes were the result of a maximum likelihood iteration procedure.

Principle and Strategy of Detection

Whenever the image of a star crossed a star mapper slit group, the pattern of the raw photon counts showed four signal peaks spaced at intervals 2:3:1 (Figures 4.9(a) and 2.1). The principle of detecting such patterns was simple: Take an integer filter consisting of zeroes and four unit values which are spaced at the same mutual distances as the four signal peaks. Fold the data with this filter, called 4-peak filter hereafter. Find peaks with a signal-to-noise ratio exceeding a certain limit, called the detection limit, in the folded data. Calculate the exact time and amplitude of this peak. Such a detected peak is called a detection, the estimated group crossing time is called transit time. The estimated amplitude is called signal amplitude.

The detection was carried out in all the Tycho data, using the sum of the simultaneously sampled counts from the B_T and V_T channels. The signal-to-noise ratio for the $B_T + V_T$ counts should usually be higher or at least equal to the signal-to-noise in a single channel. In the following we will call the B_T count rate data N_{k1}, the V_T data N_{k2}, and the $B_T + V_T$ data N_{k3}. The 4-peak folded counts in B_T will be denoted P_{k1}, in V_T P_{k2}, and in $B_T + V_T$ P_{k3}.

In the real Tycho data a simple filtering and detection scheme would have encountered several difficulties:

- the spacing of the signal peaks depended on the scanning velocity, the real star mapper geometry and the response function of the single slits;
- the method of folding the raw photon counts (additive or multiplicative) would have resulted in problems with side lobes (additive) or a loss of faint transits (multiplicative, see Volume 3, Sections 6.3–6.5);
- the expected Poisson distribution of photon count rates was disturbed by spikes (spikes are non-stellar signals not expected before launch, whose origin is inside the satellite and usually of one sample width, see Chapter 2);

- the background was highly variable and sometimes did not follow a Poisson distribution. This effect appeared regularly for background values larger than some 200 counts per sample (see Chapter 2).

The methods described below were developed to circumvent or at least limit these difficulties.

Motivation of the Non-linear 4-Peak Filter

The need of a new filtering method arose during the investigation of the first real satellite data. It immediately turned out that there were spikes in the data, i.e. photon count rates considerably higher than expected from Poisson statistics, originating from isolated events in the telescope optics. About 0.1 per cent of the data were spikes, regardless of the height of the background. Figures 4.9(a) and 2.1 show extracted parts of photon count rates in $B_T + V_T$, containing signal peaks due to a star and some spikes. Because the photon count rates were disturbed by spikes, a linear additive filtering (which in principle means just adding four samples) would have disturbed a considerable quantity of signals. Signal amplitudes would have been increased and the estimated transit times would have been disturbed.

The non-linear filtering method was at first developed to derive quantitative numbers for the spike distribution and amplitudes. Thus the filter was originally designed to do the opposite of what would be done in the normal data reduction: finding spikes, and suppressing signals from stars. However it was obvious that this logic could afterwards be turned around, i.e. that such a filter could analogously isolate star transits from spikes and, in addition, free star transits from the influence of spikes, side lobes and multiple ('parasitic') star transits.

The basic idea underlying the non-linear 4-peak filter is described by the following sentence: When the linear 4-peak filter is applied to the raw photon counts in $B_T + V_T$ (N_{k3}), the counts at the four samples p_1, p_2, p_3, and p_4 picked by the 4-peak filter should not differ more than given by Poisson statistics, loosely speaking:

$$N_{k3}(p_1) \pm m\sigma = N_{k3}(p_2) \pm m\sigma = N_{k3}(p_3) \pm m\sigma = N_{k3}(p_4) \pm m\sigma \qquad [4.1]$$

where m is a parameter defining the allowed deviation of the count rates and σ is the scatter of the count rates expected from pure Poisson statistics (both quantities will be defined precisely in the following). Equation 4.1 is also true when the filter passes the four signal peaks of a star, because each filter point p_k corresponds to the same position on the single-slit response function (see below). Equation 4.1 is not true if the filter touches upon only one signal peak (thus producing a side lobe) or when a spike disturbes the data. Any violation of the assumption made in Equation 4.1 thus indicates the occurrence of spikes, parasitic star transits and side lobes above a certain limit.

Characterisation of the Non-linear Filter

The mechanism of the non-linear 4-peak filter is as follows: Before combining $N_{k3}(p_1)$, $N_{k3}(p_2)$, $N_{k3}(p_3)$, and $N_{k3}(p_4)$ it is checked that none of them are outliers (i.e. spikes or side lobes). None must exceed a certain limit *uplim*, where *uplim* is derived from the lowest of the four count rates, $N_{k3}(p_{\min})$, which is supposed to be the least disturbed. The value of *uplim* sets the limit of the non-linearity of the filter:

$$uplim = N_{k3}(p_{\min}) + m\sqrt{N_{k3}(p_{\min})} \qquad [4.2]$$

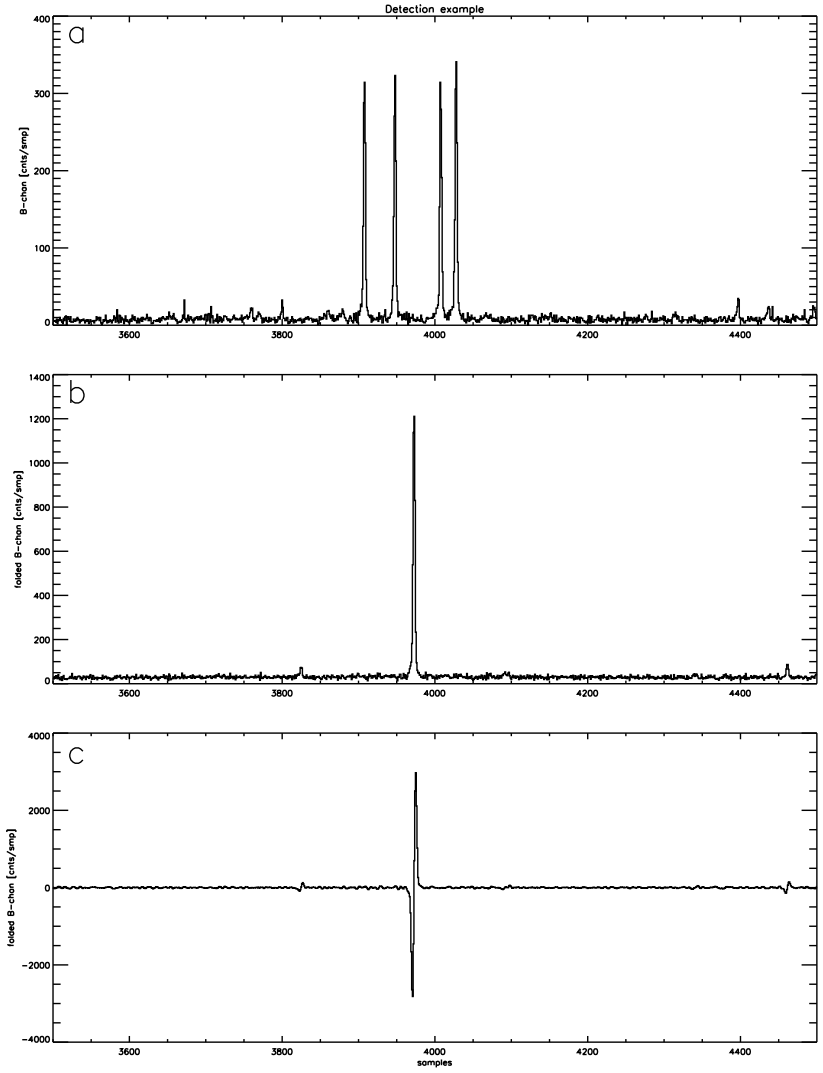

Figure 4.9. *The three main steps of the detection process. (a) The summed raw count rates (N_{k3}) of the B_T and V_T channels for 1000 telemetry samples (approximately 4.7 arcmin of scan). This stretch of data contains the transit of a bright star. (b) The same stretch of data folded by the non-linear filtering technique, resulting in P_{k3}. (c) P_{k3} folded with the 7-point transit filter, resulting in Q_{k3}.*

If the lowest value $N_{k3}(p_{\min})$ is less than the background value in $B_T + V_T$, *uplim* is derived from the background instead of $N_{k3}(p_{\min})$.

If one of the four count rates exceeds *uplim*, it is set to the mean of the other three values. But if one or more of the other three count rates are disturbed by spikes or lobes from parasitic stars, this mean value may also exceed *uplim*. In this case the count rate is set to *uplim*. When replaced, the new values $N_{k3}(p_k)$ are thus given by:

$$N_{k3}(p_k) = \min\left(1/3\sum_{\substack{i=1\\ i\neq k}}^{4} N_{k3}(p_i), uplim\right) \quad \text{for } k = 1\ldots4 \qquad [4.3]$$

These new $N_{k3}(p_k)$ are then combined to form the folded P_{k3}. Thus spikes and side lobes are bound to an amplitude which lies in a range allowed by Poisson statistics.

This non-linear filter used by detection uses the knowledge of the expected relative height of the four peaks of the signal. The difference of this filter to a purely additive (linear) filter is determined by Poisson statistics. Applying this filter to the data makes the implicit assumption, that the lowest count rates are not disturbed by any cause. Note that the filter needs a precise value for the background in order to work properly. As described in Section 4.2, either the format background or the local background was computed during the preparation of the data. The background which was actually used in the detection and estimation procedure is called the 'actual background', *b*, in the following. In general it is the format background.

The non-linear 4-peak filter may cause a bias of the estimated amplitudes, because count rates are never raised but sometimes lowered. This bias is very small and is of the order of a tenth of the overall amplitude error. On the other hand, a linear filter would cause an unknown bias towards higher amplitudes due to the presence of arbitrarily high spikes and disturbing parasitic stars.

The non-linearity of the filter starts only above a certain limit (given by *m* in Equation 4.1) in order to avoid disturbances or a bias while detecting very faint transits. This limit had to satisfy two requirements:

- an erroneous spike due to pure Poisson noise should be obtained with a probability of some 1/6000 (1 erroneous spike per format);
- transits fainter than $\simeq$10.5 mag should not be affected by the non-linearity of the filter; they were to be treated with a purely linear filter.

The appropriate value of *m* was found to be 4.5, i.e. the non-linearity started $4.5 \times \sigma$ above the expected count rate.

Practical Realisation

Despite the theoretical simplicity of the 4-point filter, the straightforward implementation of the procedure was limited by computing-time constraints. Since the filtering had to be carried out on the complete Tycho data stream of approximately 7×10^{10} data points, the practical implementation was split into two main intermediate steps. These steps introduced some arbitrary levels of significance for star transits, where the limits were chosen in a way to secure that only very few transits were lost compared to the straightforward implementation. Using this two-step procedure it was possible to decide by rather crude (and fast) criteria whether to really search for and estimate a detection or not:

(1) signal-to-noise criterion: a rough signal-to-noise ratio was computed using four neighbouring count rates around each of the filter points, p_k, of the linear 4-peak filter. If it exceeded a provisional detection limit of $0.8 \times F_{\rm det}$ (where $F_{\rm det}$ is the actual detection limit applied afterwards), i.e. only if at this place a detection could possibly be found, these data were submitted to the second step;

(2) the non-linear filtering: the second step consisted of the non-linear filtering as described above.

Searching Transits

Transits were searched in the P_{k3} stream by folding P_{k3} with the integer transit filter: [1, 1, 1, 0, -1, -1, -1] spaced over seven samples, giving Q_{k3}:

$$Q_{k3}(i) = P_{k3}(i-1) + P_{k3}(i-2) + P_{k3}(i-3) - (P_{k3}(i+1) + P_{k3}(i+2) + P_{k3}(i+3)) \quad [4.4]$$

where i is the number of a sample. A change in Q_{k3} (Figure 4.9(c)) from negative to positive values indicates a transit. A transit was accepted (i.e. turned into a detection), if its signal-to-noise ratio was above the detection limit F_{det}. But before this could be done, the rough signal-to-noise ratio had to be replaced by a more accurate evaluation.

The rough signal-to-noise ratio was computed by using four points centred on the possible transit position. If it was above 50 per cent of F_{det}, the transit time was computed by inverse interpolation in Q_{k3}. At this transit time a least-squares estimation of the $B_T + V_T$ signal amplitude a (see Section 4.4) was carried out. Taking the resulting signal amplitude a and the actual background b, a more accurate signal-to-noise ratio $F_{B_T+V_T}$ was calculated as:

$$F_{B_T+V_T} = a \Big/ \sqrt{0.15(a+b)} \quad [4.5]$$

where the factor 0.15 is a trivial consequence of the normalisation of a and b. If $F_{B_T+V_T}$ was above the detection limit F_{det}, the transit was accepted. The detection limit used inside a Predicted Group Crossing interval ($F_{\text{det}} = 1.5$) was different from that outside ($F_{\text{det}} = 1.8$). These limits were chosen in order not to get too many false detections and not to lose many true transits. For stars with $V_T = 11.5$ about half of the true transits were found, the other half being censored. One accidental false detection was found per Predicted Group Crossing interval. Due to the higher detection limit outside these intervals a smaller amount of false detections was found there.

The detection of a transit, as described above, included the estimation of the $B_T + V_T$ transit time and amplitude. The $B_T + V_T$ amplitude obtained from P_{k3} could be slightly underestimated due to the not exactly spaced 4-peak filter, as explained below. The transit time estimate remained unaffected, because the barycentre of the 4-peak filter still corresponded to the barycentre of the signal peaks, and any deformation of the folded peak remained symmetric. For amplitude estimations in B_T and V_T, N_{k1} and N_{k2} were folded separately, with the inexact spacing of the 4-peak filter taken into account (see below). Thus B_T and V_T amplitudes were not disturbed by this effect.

Folding of the B and V Photon Counts

After the search for transits in N_{k3} and the $B_T + V_T$ transit time estimation, the folding of N_{k1} and N_{k2} was carried out. No further signal-to-noise check was done, because a transit was supposed to be there. Once the transit time was fixed, the integer spacing of the 4-peak filter could be taken care of. Due to the undersampling of the Tycho data the gradient of the count rates was rather steep, especially for bright stars. Thus in extreme cases the non-linear filter could have found count rates exceeding the *uplim* criterion just because of the integer filter positions. To prevent these cases the following procedure was implemented: using the (known) transit time, the actual scanning velocity and the precise slit group geometry for each filter point, the approximate position on the corresponding slit, i.e. the value of the single-slit response function at the filter point positions, was computed. In principle, the folding was done in the same way

as the N_{k3} folding. But while applying the non-linear filter, the four points under consideration were normalized to the same value of the single-slit response function. Before combining the points, the normalisation had to be reversed, of course, in order to preserve the original shape of the signal.

4.4. Estimation of Transit Parameters

Transit Time Estimation

B_T and V_T transit times were estimated separately by using the same integer filter method as for $B_T + V_T$ transit times, but operating on P_{k1} and P_{k2}, which result from the non-linear filtering in B_T and in V_T. It was shown by Yoshizawa *et al.* (1985) that a maximum cross-correlation filter method for the estimation of transit times is not more accurate than this simple integer filter method.

Signal Amplitude Estimation

Yoshizawa *et al.* (1985) showed that for the signal amplitude estimation a maximum-likelihood method is superior to a least-squares fit. Thus, the B_T and V_T amplitudes were evaluated using a maximum likelihood iteration, the starting values being the results of a least-squares fit. Input required for the amplitude estimation were the actual background, the transit time estimate, and the relevant response function (vertical/inclined slit group, preceding/following field of view, upper/lower half of slits, B_T / V_T channel) and folded photon counts in the B_T and V_T channels.

Preparing the single-slit response function: For the fitting procedure a folded slit response function F_p had to be prepared to correspond to the 4-peak folded data. This means the following: Assume the estimated transit time p corresponds to the barycentre of the four slit signals. Knowing the actual scanning velocity and the slit geometry, the distances of the four slit signals (in units of samples) and the approximate position of each sampling point on the corresponding slits can be computed. Thus it is possible to reconstruct the real signal peak at the transit time by adding the relevant values of the single-slit response functions. The shape of the resulting folded slit response function F_p depends on the fraction of p (p in units of samples) and on the actual scan speed.

Least-squares estimation for amplitudes as starting value: The result of a least-squares fit of the folded response function to 10 points of the 4-peak folded data was used as the starting value for the final maximum likelihood amplitude estimation. The least-squares fit determined only the signal amplitude a, keeping the previously determined background and transit time fixed. The minimum of the sum $\sum_i (aF_p(i) + 4b - P_k(i))^2$ over 10 data points was determined as:

$$a = \sum_i (P_k(i) - 4b)F_p(i) \Big/ \sum_i F_p^2(i) \qquad [4.6]$$

where i denotes the indices of the 10 data points.

Whenever this amplitude estimate resulted in a zero or negative value, the shape of the signal did not fit the expected one. In this case the transit was either disturbed or it was a false detection. If this occurred while estimating the $B_T + V_T$ amplitude during the

search for transits, these transits were not accepted. If it occurred while estimating B_T or V_T amplitudes, the transit was flagged 'bad' in the corresponding channel, and no further estimation was done in this channel.

Maximum likelihood estimation of final amplitude: The final maximum likelihood amplitude estimation was an iterative method, giving the most probable amplitude a, assuming the previously determined transit time p and the background b to be correct. Photon count rates were supposed to be Poisson distributed. The result was obtained by a Newton-Raphson iteration process (Silvey 1970):

$$a_n = a_{n-1} + \sum_i (P_k(i)/P_{n-1}(i) - 1)F_p(i) \Big/ \sum_i F_p^2(i)/P_{n-1}(i) \qquad [4.7]$$

where i denotes the indices of the 10 data points, P_{n-1} is the theoretical P_k, i.e. $P_{n-1} = 4b + a_{n-1}F_p(i)$.

The iteration was stopped as soon as a_n and a_{n-1} differed by less than 0.1 per cent. Again, if the amplitude iteration did not converge, the transit was flagged 'bad' in the corresponding channel.

4.5. Verification Methods

The verification of the software used for the prediction and detection tasks, and of their output data streams, was a continuous process. It started well before the launch of the satellite and ended only after a first preliminary Tycho output catalogue had been produced by the astrometry task. It had to be an iterative procedure, primarily because of the gradual development of the software, but also because of the iterative nature of the Tycho data reduction chain itself. Verification efforts before launch included the trial processing of simulated satellite data provided by NDAC. After launch, the verification processes were interleaved with the calibration of the instrument, using actual mission data.

The prediction and detection processes are so intimately connected that much of the verification of prediction could only be done with the output of detection in hand. Thus the verification of either of these processes was at no stage independent from the other one. Verification of prediction was furthermore made difficult by the fact that prediction had to be accurate to the milliarcsec level in order not to degrade the final Tycho astrometry. But, due to the low precision of the real-time satellite attitude and of the Tycho Input Catalogue, the actual agreement between predicted and observed group crossings could initially be tested to the level of a few tenths of an arcsec only. This situation only gradually improved with the advent of on-ground attitude, of more precise star mapper calibrations, and finally with the first preliminary Hipparcos output catalogues. More details can be found in Chapters 7 and 10.

The remainder of this section briefly outlines some of the early verification steps, starting with different aspects of the prediction software, and proceeding to the detection task later on.

Satellite Orbit Comparison

A correct interpretation of the satellite ephemerides (which were provided on each of the telemetry tapes, see Section 4.1), was crucial to the computation of stellar aberration and (for the solar system objects) of parallax. Shortly after launch, the European Space Operations Centre provided a sample orbit file to the Data Reduction Consortia, along with a numerically tabulated interpretation of the Chebychev polynomials into satellite position vectors and velocity vectors, in the different relevant coordinate systems, for some specific instants of time. These data were compared with the output of the relevant subroutines in the operational prediction software.

Apparent Positions Comparison

The computation of apparent star positions from stellar data (mean position, proper motion and parallax), solar system ephemerides and satellite position and velocity vectors was verified by comparisons among the different Data Reduction Consortia and the European Space Operations Centre. In 1990, a set of *ad hoc* comparison inputs was produced by NDAC (the Royal Greenwich Observatory group), and successfully passed by all parties. In 1992, another comparison was performed by extracting apparent positions actually computed during the operational runs of the data reduction programs of FAST, NDAC and TDAC. For this purpose, a few particular stars and a few hours of mission time were selected at random. The analysis of the resulting sample outputs showed agreement on the 0.1 mas level among all four participating parties.

Star Selection for Prediction

The most complicated part of the prediction software concerned the Tycho Input Catalogue access, i.e. the selection of those stars that may produce group crossings within a certain small time interval. If this procedure had been imperfect, the Tycho data reduction chain would have lost part of the observations. This could not be fully verified before launch. It was partly tested by, for example, manually constructing the run of the star mapper slits over a part of the Input Catalogue sky, and comparing with the actual Predicted Group Crossing data stream. Final verification came much later by checking (over a significant part of the mission) that essentially none of the actually observed transits of bright stars (by the detection process) remained without corresponding Predicted Group Crossing data. Still another check was to look for 'holes' in the sky distribution of all stars having Predicted Group Crossing records. The only 'holes' found after a few months of mission corresponded to 'holes' in the Input Catalogue, which in turn could be verified as real by comparison with the printed photographic sky surveys. Another verification of the completeness of prediction came from the star recognition process described in Chapter 5. The few bright ($V < 8.5$ mag) stars missing from the Tycho Input Catalogue Revision were thoroughly investigated. None were due to missing Predicted Group Crossings.

'Cloud Plots'

Plots of the individual transit time differences 'detected minus predicted' versus the z coordinate along the slits were referred to as 'cloud plots'. Figure 4.10 shows a few

examples. Each of the tiny × symbols corresponds to one detection inside a Predicted Group Crossing interval. The height of the plots is slightly less than the full width of the interval. Units on the vertical scale are 1/600 s, the time interval of one Tycho photon counts sample. On the sky, this corresponds to 0.281 arcsec for the vertical slits (right half of the figure), and to 0.199 arcsec for the inclined slits (left half), at the nominal scan speed of 168.75 arcsec s^{-1}. The detections in the dense 'clouds' are those that are really related to the predicted star transits, while the homogeneous background of detections on either side of the clouds are due to photon noise peaks, and to unrelated stars crossing the slits at about the same time.

The cloud plots were one of the major diagnostic tools for prediction and detection. The width of the clouds are indicative of the combined precision of the attitude and input catalogue. Any offset of the clouds from the centre of the Predicted Group Crossing intervals, on the other hand, shows errors or inconsistencies in the interpretation of the attitude and/or instrument geometry. The examples in Figure 4.10 were chosen to illustrate some of the early improvements achieved at the start of the mission. They do not all belong to the mass production phase of prediction and detection, but to the extensive trial runs with real satellite data that were performed during 1990.

The two plots at the top contain the first results on a stretch of data collected before 23 January 1990. The inclination and offset of the clouds are caused by the unexpected rotation of the grid assembly relative to its nominal position in the focal plane. The 'jump' of the inclined slits at $z = 0$ is also caused by this effect. The rms width of the clouds is about 0.9 arcsec, which is mainly due to the imprecise real-time attitude at this early stage of the mission. This state of affairs is clearly unsatisfactory: some real star transits are being missed at the edges of the Predicted Group Crossing intervals, especially at large negative z coordinates (lower left in the plots).

The two plots in the centre show analogous data, but with the grid rotation (and a systematic attitude offset) taken into account in the prediction process. This removed the systematic offsets, but could not improve the random errors of the attitude, i.e. the width of the clouds.

The two plots at the bottom show data collected after 23 January 1990, when the grid rotation had been implemented into the on-board instrument model. The improvement is evident. It is due to the improved real-time satellite attitude. The rms width of the clouds is about 0.6 arcsec for the inclined slits, and 0.4 arcsec for the vertical slits. It is apparent that the instrument model used in the prediction process for these data caused a slight overcorrection of the 'jump' in the inclined slits at $z = 0$. However, the defect was sufficiently small not to disturb the subsequent Tycho data reduction steps.

Standard Quality Control in Detection

During the mass processing of detection a set of standard tools were used in order to monitor the quality of the input and output data. These standard tools used statistical data collected during the routine processing. A total of about 0.5 Gigabytes of such data were produced and regularly inspected. Besides the 'cloud plots' described previously, a set of 10 different sorts of plots were provided by the so-called 'Diagnostic Package' reading the statistical data sets. Some of these plots are presented and described hereafter for one arbitrarily selected stretch of satellite observations. The plots show the typical properties of the input and output of the detection and estimation process for reprocessing data.

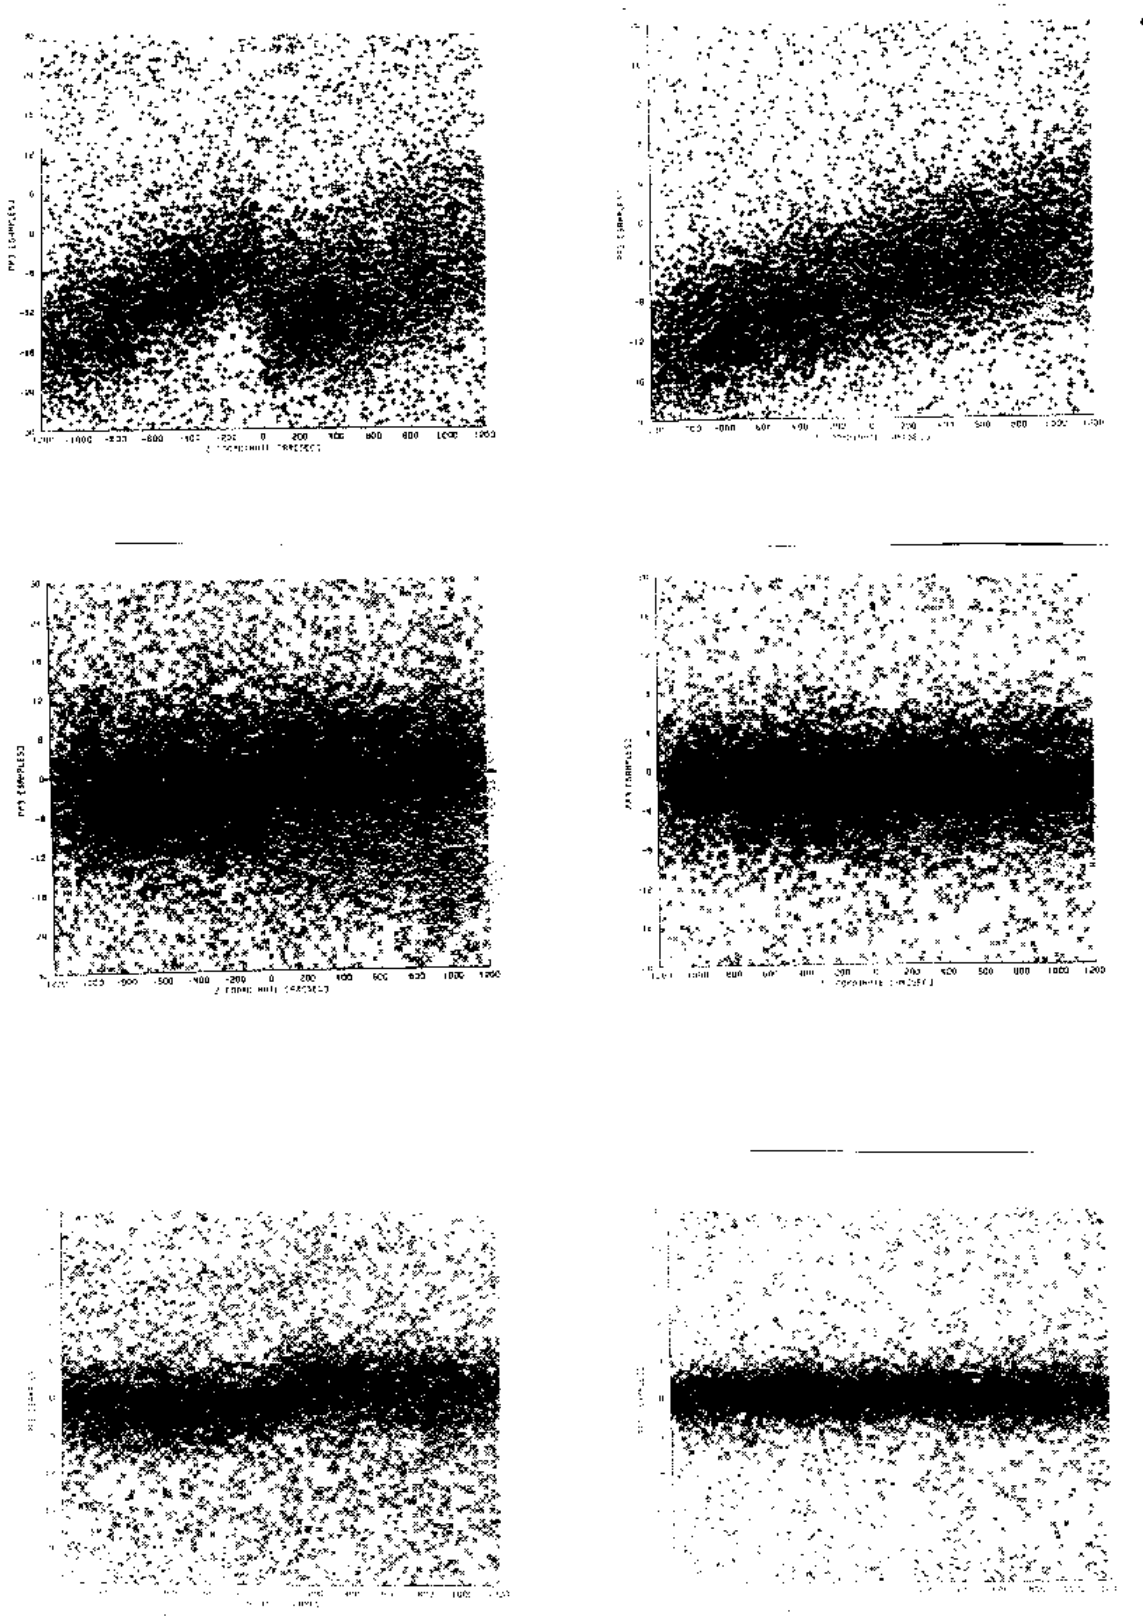

Figure 4.10. *Sample 'cloud plots' from the early stages of trial processing of prediction and detection. The height of the plots corresponds to 12 arcsec on the sky, i.e. the width of the Predicted Group Crossing intervals. The horizontal axis spans the 40 arcmin length of the star mapper slits. Details are described in the text. The improvement from top to bottom is evident.*

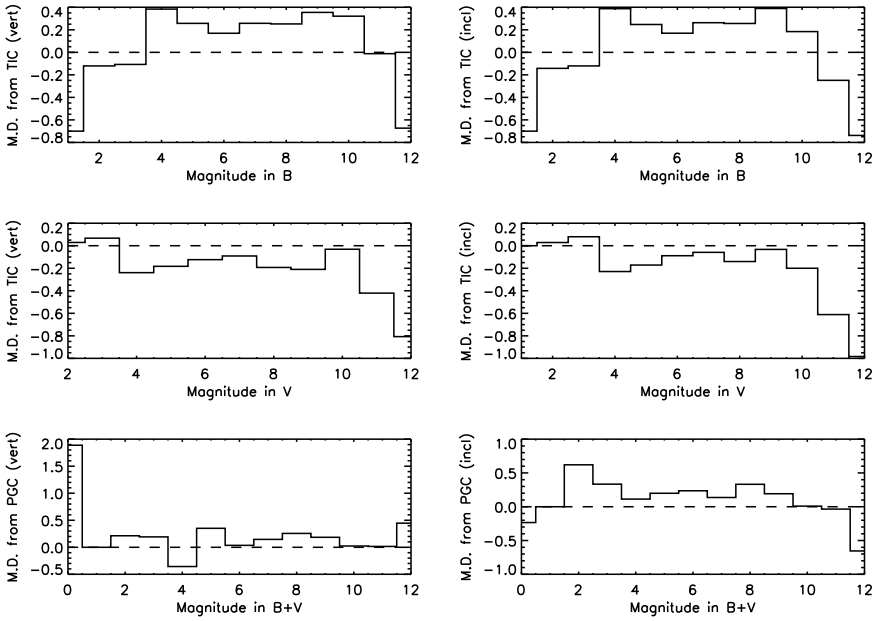

Figure 4.11. *The mean difference between the observed and the Tycho Input Catalogue magnitudes (derived from a crude calibration) is shown for different magnitude bins in B_T (top) and V_T (centre), for the vertical (left) and inclined (right) slit groups. The roughly constant values throughout the centre of the plots indicate linearity of the estimation process. The number of bright stars is very low. Thus the strong variations at extreme left are due to small-number statistics. At extreme right, i.e. at very faint magnitudes, the effect of the censoring causes a drop of the plotted curves, but again small-number statistics modifies the effect. The lower panels show the mean difference between the predicted and the measured transit times versus $B_T + V_T$ magnitude in the vertical (left) and the inclined (right) slit group, respectively, in units of samples (corresponding to 0.281 arcsec on the sky).*

Mean deviation from input catalogue magnitudes: The signal amplitudes derived by the estimation process were converted to magnitudes, using a crude calibration formula provided by the photometry task. For a linear photometric system, the mean of the differences between the measured magnitudes and the Tycho Input Catalogue magnitudes should not depend on the magnitude (the mean value itself is of no importance; it is just the error of the crude calibration zero point used). Since the photometric calibration did not use a magnitude-dependent term (see Table 8.2), any magnitude dependence (i.e. non-linearity) in estimation would be reflected in the finally reduced magnitudes. Plots like the one shown in Figure 4.11 were used to check the linearity (i.e. constancy of the offset versus magnitude).

Predicted minus observed transit times: The detection process depended very much on the accuracy of the predicted group crossings. Thus the monitoring of the quality of prediction was vital for the whole data reduction. One kind of quality control was the production of the 'cloud plots', see Figure 4.10. However, the production of these plots was quite time-consuming and thus impractical for daily use. So-called 'PDI-plots' (Figure 4.12) were used as a quick-look utility on the same parameters, namely the difference between the measured and the predicted crossing, but without resolution along the z axis.

Goodness-of-fit: During the estimation process the goodness-of-fit (chi-square per degree of freedom, χ^2/f) was determined for every detected transit. In an ideal instrument, the frequency distribution of the goodness-of-fit should follow a normalized

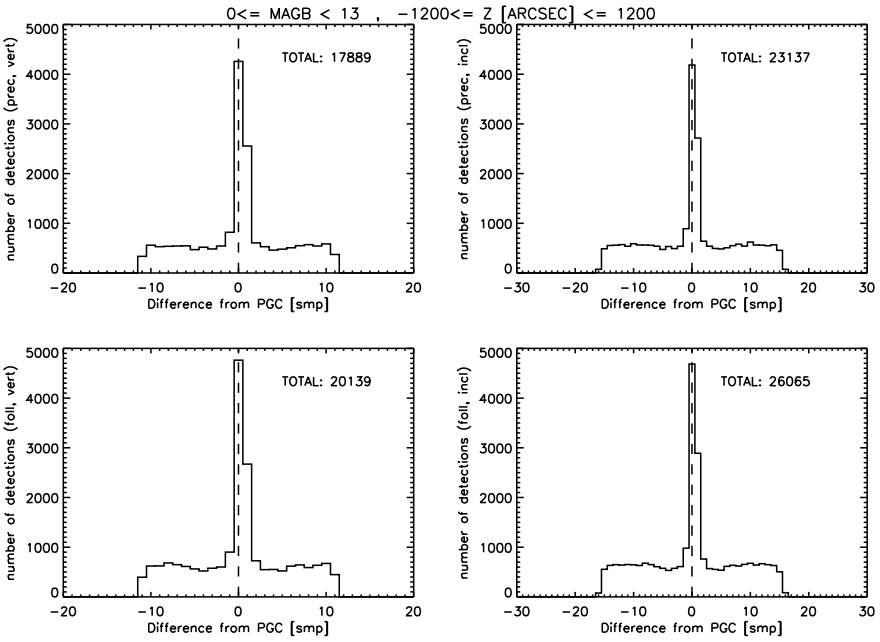

Figure 4.12. *Histograms of the differences between the observed $B_T + V_T$ transit times and the Predicted Group Crossing epochs are shown for the vertical (left) and inclined (right) slit groups, and for the preceding (top) and following (bottom) field of view. The width of the plots corresponds roughly to the Predicted Group Crossing interval. The peaks in the centres are due to the predicted star transits, while the broad plateaux are due to unrelated ('parasitic') stars and photon noise ('false') detections. Only in the reprocessing are the peaks as narrow as shown here. In the main processing stage their total width was of the order of ±4 samples for the vertical slits.*

χ^2/f distribution with a peak around $\chi^2/f = 1$. In a more realistic case, the distribution of the mean goodness-of-fit versus magnitude should show a smooth behaviour. The actual distribution for the Tycho data depended on a number of parameters such as the background and the magnitude distribution of the stars contained in a specific data set. The actual values derived in the estimation process showed a strong dependence on the magnitude. At faint magnitudes where the differences between observed and modelled count rates were dominated by photon noise, the actual goodness-of-fit roughly corresponded to the ideal expectations. The strong increase towards bright magnitudes (Figure 4.13), is due to the fact that there the photon noise becomes negligible compared to other error sources. The increase is due to, for example, small errors of the model response functions, the undersampling of the response functions, and other effects, which are all more disturbing for bright than for faint stars. For stars brighter than about $V_T = 4.5$ mag or $B_T = 5$ mag, the goodness-of-fit usually reached and even exceeded 12.7, the maximum value which could be stored in the range provided by the output formats. This is the reason for the high plateau in the top and centre panels of Figure 4.13, and for the narrow peak in the two lower panels.

Signal-to-noise ratio: The signal-to-noise ratio of detected transits should obviously depend on the magnitude of the observed stars. This is confirmed by the upper part of Figure 4.14 (the levelling-off of the curve for stars brighter than about 7.5 mag is due to a cutoff imposed by the data format, as in the case of Figure 4.13). However, it should not strongly depend on the background, at least within the background range used by photometry. This requirement was monitored by plots like the ones shown in the lower part of Figure 4.14.

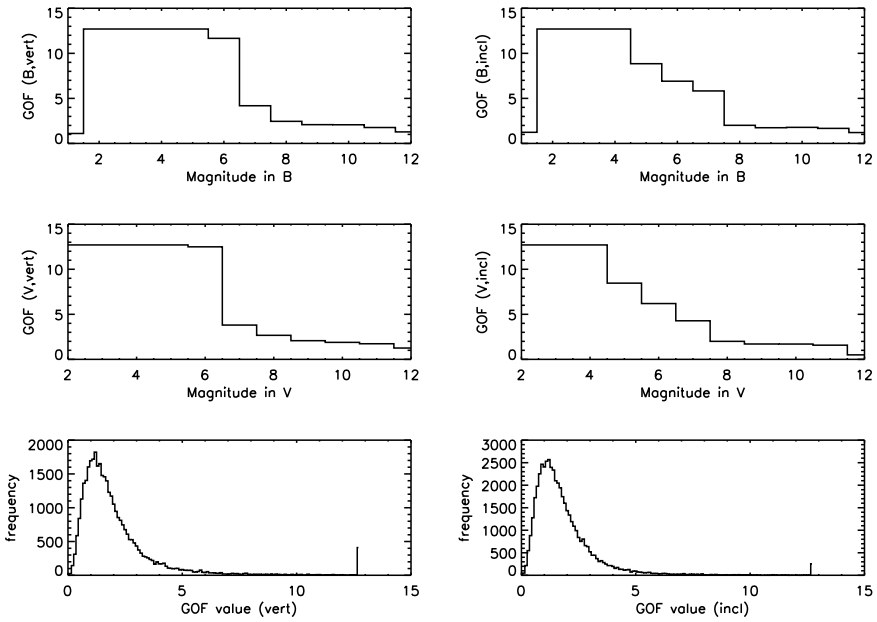

Figure 4.13. *The mean values of the goodness-of-fit (chi-square per degree of freedom) in B_T (upper) and in V_T (middle) is shown versus magnitude for the vertical (left) and inclined (right) slit groups. The excess towards brighter magnitudes results from the very high sensitivity of the fits to the steep slopes of the response function. The overall distribution of the goodness-of-fit is shown in the lower two panels, again for the vertical (left) and inclined (right) slit groups. Details are explained in the text.*

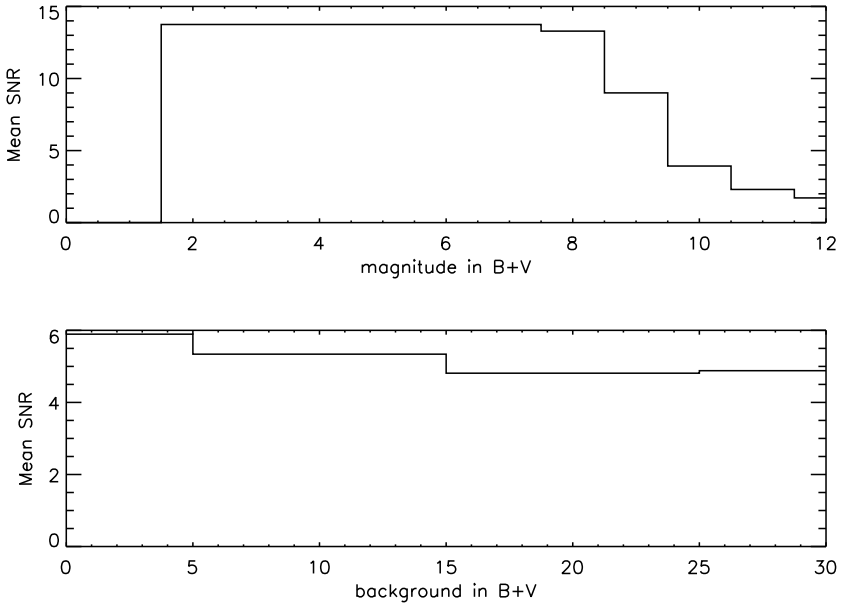

Figure 4.14. *The mean signal-to-noise ratio versus B_T + V_T magnitude is shown in the upper graph. In the lower graph, the mean signal-to-noise ratio versus background is shown.*

Background: The background observed in the Tycho detectors was one of the worst consequences of the highly elliptical orbit of the revised Hipparcos mission. Usability of the raw data was almost solely a function of the background. Thus the distribution of the background level was one of the most important quantities to be monitored. Within the 'Diagnostic Package' there were two kinds of plots which gave an instant overview of this parameter. In addition, there was a quite extensive external toolset for off-line visualising and investigation of the background. Some of the results of this toolset are shown in Figures 4.3 to 4.7.

Step-by-step reduction: Another very important diagnostic tool was a special version of the complete detection and estimation (plus rough photometric reduction) software package providing a manually controlled step-by-step treatment of the data for a single, freely selectable star, for a given stretch of satellite observations. It provided the possibility of visualising details of each processing step. This tool proved invaluable in finding software errors, investigating data anomalies and interpreting problem stars.

U. Bastian, A. Wicenec