4. OVERVIEW OF THE DATA ANALYSIS

This chapter provides a general overview of the data analysis leading from the raw satellite data to the Hipparcos Catalogue and associated annexes—the details of the various processes are presented in the following chapters. The basic concept for the data analysis had already been established at an early stage in the mission planning, and the reductions were subsequently carried out in a series of sequential processes, each substantially decreasing the corresponding data volume. Certain aspects of the analysis were in practice dictated by the relatively limited computer resources at the time of development and earliest implementations. Much effort was devoted during the preparatory phases to creating reduction software that was both fast and highly reliable. The practical organisation of the treatment of the data in both consortia, FAST and NDAC, is also summarised.

4.1. Main Stages of the Data Reduction

Input Data Stream

The input data stream consisted of the scientific data from the satellite, auxiliary data (principally the satellite orbital parameters) and various satellite 'housekeeping' data. These data streams were prepared by ESOC from the original telemetry data, and put on 9-track, 6250 bpi tapes (see Volume 2, Chapter 9) at a rate of just under 1 tape per day. Most of the data was provided in a highly compressed form, with some 1350 tapes in total delivered by ESOC to each consortium. The total amount of compressed science data relevant for the construction of the Hipparcos Catalogue was around 70 Gigabytes.

Reduction Processes

A 'three-step' reduction scheme was first proposed by L. Lindegren, of Lund Observatory (Sweden), during the feasibility study of the space astrometry mission. It consisted of three major steps: the great-circle reduction, the sphere solution and the determination of the astrometric parameters. Although a direct, global solution of the mission data would have been possible in principal, it was totally impossible in practice. The three-step method allowed the overall analysis problem to be decomposed into three sequential steps, allowing the processing of the satellite data to be considered as feasible. It led to a marginal degradation of accuracy compared with a theoretically optimum reduction system, and the same general approach was adopted by both consortia. However, steps 2 and 3 were combined to one step by NDAC in the final implementation, which also comprised the processes for star mapper data reduction, attitude reconstruction and image dissector tube data reduction, plus a wide range of off-line processes (see below).

In practice, the processing did not follow an exclusively sequential structure—thus, for example, results from the sphere solution were used to iteratively improve the attitude reconstruction and great-circle reduction, by feeding back improved stellar positions into the reduction process, while results from the great-circle reduction were used as a quality control of the attitude reconstruction. Several other processes ran largely in parallel with the main data processing, but with a sometimes complex interaction with other processes:

- the double star processing (Chapter 13) interacted at many points with main reduction processes: results from the star mapper processing provided precise starting coordinates for double stars with separations larger than 1.5 arcsec, image dissector tube data processing provided the input data for the double star processing, the great-circle reduction provided relative reference positions for the measurements, and the sphere solution results allowed these to be transformed into absolute positions;
- the photometric reductions (Chapter 14) proceeded rather independently of the astrometric tasks, although the photometric results were used in all three main processes, in particular providing colour determinations obtained with the star mapper data;
- comparisons between various calibration parameters obtained at various stages of the data processing led to a better understanding of the instrument and thus of the reliability of the reduced data. These aspects are presented in Chapter 8 and 10;
- special treatment was required for the processing of minor planets and planetary satellites (Chapter 15), from the first processing to the presentation of the final results;
- at all stages of the reductions, a catalogue of stellar parameters (positional and photometric) was used, and was regularly updated using intermediate mission results.

The rest of this section provides some more details about the reduction processes and references to the chapters in the current volume where full descriptions can be found. The processes are here divided into the main reduction chain, the parallel processes, and other aspects.

Main Reduction Chain

Part A. Processing of photon counts (Chapters 5 to 7): The aim of the first reduction processes was to derive from the image dissector tube counts, in combination with the star mapper data (and to some extent gyro data), the phases of the modulated signals at a given reference time during the star transits, and with respect to a chosen reference line on the grid. In addition, the satellite attitude angles at the chosen reference time were needed for further processing, as well as for the first processing of the image dissector tube data. The processing proceeded as follows:

(1) preliminary investigation of the satellite attitude to provide input data required for the star mapper processing, processing of the star mapper data to transit times (Chapter 6);

(2) using catalogue positions and star mapper transit times, the attitude for each observational frame with image dissector tube data was reconstructed with an accuracy of better than 0.1 arcsec (Chapter 7);

(3) using the reconstructed attitude results, the image dissector tube data were reduced to provide phase and amplitude information for the astrometry and photometry respectively (Chapter 5).

The output of the latter two processes formed the input for Part B, the great-circle reductions (Chapter 9). Parallel processes in Part A were the star mapper and image dissector tube photometric reductions, catalogue improvements, and a range of calibrations. Data were also prepared for the double star processing. The total amount of data at this stage was, on average for each observation frame of 2.13 s, 4.5 transit times, their accuracies and amplitudes, the three satellite attitude angles with accuracies, as well as full timing information. The reduction in data volume from the original photon counts was about a factor of 100. On average, about 50 000 transit times were obtained per orbit (these typically constituting one 'reference great circle').

Part B. The great-circle reductions (Chapter 9): The task of the great-circle reductions was to combine the 50 000 transit times obtained over one orbital period into abscissae along a reference great circle, for the 2 000 different stars observed during that period. This was obtained through a further refinement of the along-scan satellite attitude, through projection onto a reference great circle using the results of process (2) described under Part A, and through calibration of the large-scale distortions of the projections on the main detector. Also calibrated was the basic angle (Chapter 10) between the two fields of view, which enabled the linking of different parts of the sky at the milliarcsec level. The monitoring of the various calibrations was used as part of the overall quality control.

Double star coordinates were generally not carried along in the great-circle reduction: the reference phase of the modulation for a double star was a combination of two signals, the result of which depended on the orientation of the modulating grid with respect to the orientation of the double system. They had to be treated independently, using, however, the scan-phase information obtained in the great-circle reduction.

The output of each reference great circle reduction was a reference pole, a preliminary zero-phase (relative to some celestial reference frame), and abscissae, and their accuracies, for some 2000 stars. This represented the data collected for single stars over an interval of between 6–9 hours.

Part C. Sphere solution and astrometric parameters (Chapter 11): The task of the sphere solution was to establish a consistent system of zero-phases for all reference great circles, and subsequently derive from the combined abscissae, corrected for the zero-points, the five astrometric parameters: the two components of position, the two components of proper motion and the parallax. In the process, effects that could not be detected at earlier levels in the data reductions were calibrated: certain harmonics in the great-circle solutions could enter due to occultation-gaps, and chromaticity corrections could only be obtained at this level.

The final astrometric parameter determination was carried out after the merging of the data from the two consortia (Chapter 17). At this stage non-linear proper motion cases were also detected and solved for.

Parallel Processes

A series of parallel processes used data from the main processing chain, and sometimes provided information to it. The two principal parallel processes were the double star processing and the photometric calibrations.

Double star processing (Chapter 13): The character of the modulated signal obtained with the main detector was such that all information on double stars had to be processed as a separate, dedicated task. Input data came from the image dissector tube data processing, the image dissector tube photometric calibrations, the great-circle calibration results and the reference great circle zero-point calibrations obtained as part of the sphere solution. Further information for the detection of double stars was provided through accumulated statistics from the image dissector tube data processing. Data from the optical transfer function (Section 5.9), describing the characteristics of the modulated signal as a function of position in the field of view and as a function of star colour, was incorporated in the reduced image dissector tube data results. These data were not only supplied for double stars known *a priori*, but for every star observed, since many unknown double stars were present in the observing programme.

The double star processing consisted of two stages: (1) recognition of stars as double (or multiple); (2) solving the double star parameters (separation, orientation, magnitude difference). By using data from the sphere solution and the great-circle reductions, the double star parameters could be entered into the final sphere solution. Checks on the reliability of this fitting were made through treating some single stars with the same processing, and comparing the results with those obtained in the normal processing.

Photometric calibrations (Chapter 14): Two types of photometric calibrations were carried out: for the star mapper data, and for the image dissector tube data. The first was of moderate accuracy, and intended to provide colour information for stars without reliable colour information in the Hipparcos Input Catalogue. This could only be done for stars brighter than approximately 10 mag. The star mapper photometry also provided a background signal relative to which the background signal in the main detector could be determined.

The image dissector tube photometry provided accurate photometric data in a broad band, eminently suitable for use as 'epoch photometry' for the detection and investigation of variable stars. It was also used as calibration information in the double star processing.

Catalogue updates (Chapters 6 and 16): The quality of the data reductions was considerably improved through the work done on updating the catalogue information: improved positions provided a more accurate attitude reconstruction, allowed a better distinction of the pointing of the instantaneous field of view, and a more reliable great-circle reduction (through removal of grid-step ambiguities). Catalogue updates were initially provided through accumulated star mapper reduction results: the attitude reconstruction was more accurate than that represented by the positions of stars in the Hipparcos Input Catalogue, and the residual transit times left after attitude reconstruction were used to correct those positions. The star mapper data also provided improved colour information as described in the previous section. Finally, the star mapper data provided improved parameters for some 1200 double stars. At later stages in the processing, improvements came from the first preliminary sphere solutions, at which point positional accuracies were reached that no longer had any negative influence on their use in the reduction processes.

Catalogue updates were also provided to the satellite operations team at ESOC for improved satellite performance.

Other Aspects

Calibrations (Chapters 8 and 10): Calibrations were carried out at almost every stage of the data reductions, and comparisons between various calibration results led to a much better understanding of the instrument and the external influences on it. In many cases calibrations were performed to a level well beyond the strict requirements of the data processing, providing information that can be of use in future space missions.

Interfaces with the Tycho Catalogue reductions: The Tycho data reductions required the attitude reconstruction from the main mission. As the Tycho reductions used the star mapper data, their reduction processes and results were incorporated in comparison exercises. The Tycho photometry was essential in establishing the reference colours for the final main mission reductions.

Comparisons: Careful and extensive comparisons were carried out throughout the data reductions, an important process which began before the start of the mission using simulated data, and which was to some extent facilitated by the common adoption of the same sequence of reduction processes by both consortia. The main comparison tasks were:

(1) star mapper reductions: detections, transit times, intensities, and error estimates;

(2) attitude reconstruction: coverage, general performance, outliers;

(3) image dissector tube data reductions: phase and amplitude estimates, results from ‘partially observed stars’, accuracies;

(4) great-circle reduction: abscissae, accuracies, slit ambiguities;

(5) sphere reconstruction: accuracies, rotations;

(6) double stars reductions: detections, parameter comparisons;

(7) photometric reductions: biases, error estimates, data rejection and flagging, barycentric corrections;

(8) variability investigations: detections, periods, zero-phases;

(9) satellite orbital parameters, and determination of celestial directions.

All these comparisons had important consequences for the mission results: they revealed numerous errors or possibilities for improvement in the reductions, stimulating the data reduction groups to produce improved and optimum results—this element of (good hearted) competition was an important factor in obtaining the final high quality of the Hipparcos Catalogue results.

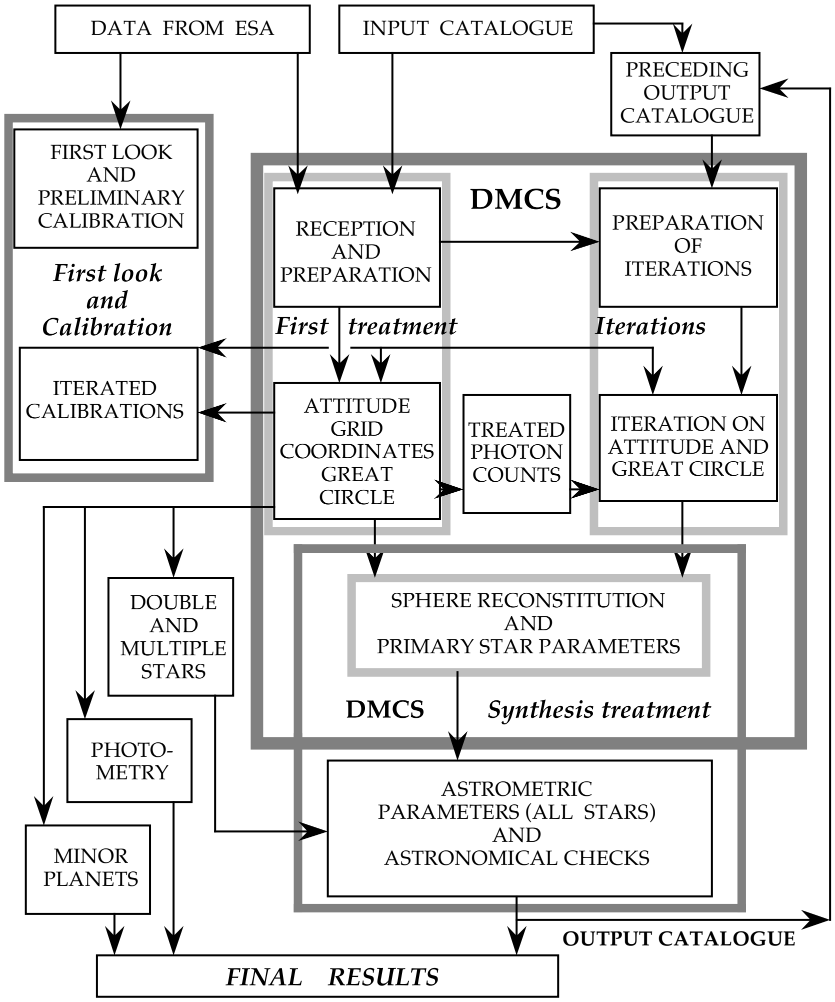

Figure 4.1. *General scheme of the interactions between tasks in the FAST Hipparcos data reduction organisation. DMCS: Data Management and Command System.*

4.2. Organisation of the Data Reductions in FAST

The data analysis by the FAST Consortium followed the principles described above, being essentially an iterative process during which intermediate quantities such as attitude and instrumental calibrations status were improved together with the five astrometric parameters (see Figure 4.1). The process was iterated about ten times, either when new observational data was included or simply when a better working star catalogue was available from earlier reductions. On several occasions, software was improved between two successive iterations, partly as a result of the comparison exercises, also thereby contributing to the general improvement of the results.

The various tasks were prepared in different scientific institutes before the launch of the satellite. They were verified using simulated data prepared in CERGA (Observatoire de la Côte d’Azur), and put in a format that was pre-defined in an architecture document and an interface document where all the interactions and data exchanges between software were foreseen. The actual exploitation of the software was carried out systematically on the CNES CDC computer, so that the software had to be transferable, irrespective of the computer used in the contributing institutes. Thus each piece of software had to be verified first in the responsible institute then, once transferred, in CNES. An ensemble of acceptance tests was set up for each software module, and acceptance was declared after the tests had been fully satisfied. An inter-institute committee, referred to as the Software Advisory Group (SWAG), which included representatives of CNES (J.L. Pieplu and C. Huc), CERGA (J. Kovalevsky and J.L. Falin) and other participating institutes (E. Canuto, CSS; R. Hering, ARI; H. Kok, Delft; and F.P. Murgolo, CSATA) played a major role in building the interfaces, in enforcing inter-institute cooperation, and in transforming independent software into a coherent data processing system.

The role of the Software Advisory Group continued after satellite launch. The software accepted on the basis of simulated data, had to be proved adequate to treat the real data. All the software modules, separately and in combination, had to be confronted with unforeseen problems or perturbations, and with the consequences of the highly eccentric orbit such as the absence of data during perigee passages, the increase of photon noise in the van Allen belts, long eclipses, and the large variability of torques with time.

It took more than one year to rewrite some of the software and to re-perform all the acceptance tests on the CNES computer using some data sets chosen in such a way that all different peculiarities of the orbit and of the data recovery were fully covered. Thus the actual systematic data treatment started only in May 1991, 1.5 years after the actual start of the mission. The automatic mass treatment organised in CNES, and controlled by the 'Data Management and Command System' was very efficient, and the lag rapidly diminished. However additional improvements, in particular in attitude determination and great-circle reduction, continued to be made for at least another year and a half.

The FAST reduction structure can be described by grouping the various tasks and related software into five categories:

(1) data reception: this consisted of transforming ground-based data (the Hipparcos Input Catalogue, and planetary ephemerides) and the satellite data transmitted by ESOC into a form directly accessible by the various software modules;

(2) 'first-look' and calibrations: every week, the first-look task in SRON, Utrecht treated one data set (data acquired during one satellite orbit) to validate the data sent by ESOC almost in real time, and to provide the first calibrations to the FAST data reductions. Later, these calibrations were improved using the results of the main data reductions;

(3) mass treatment: this consisted of transforming the raw data into star abscissae on a reference great circle, carried out data set by data set. Two different chains were used: one for the first treatment of the data, the second for the iterations;

(4) synthesis: all the available results on data sets, whether obtained from the first or iterated treatments, were combined to determine a consistent mesh of reference great circles with their origins, followed by the astrometric parameters of each star;

(5) off-line tasks: these included photometry, double and multiple star treatment, minor planet reductions, and the improvement of calibrations.

Each of these tasks, and their interactions, is described hereafter.

Data Management and Command System

The decision to perform all the data treatment on a single large centralised computer, and to do it as automatically as possible, required the setting up of a software system which would be able, using simple commands, to process any subsystem for any data set, and to access all intermediate or final files at any time.

Processing all the data, over the whole mission, demanded the correct and coherent execution of more than 30 000 jobs, accessing 80 classes of files, comprising a total of more than 120 000 files. In order to minimise the consequences of mistakes, to control the progress of the processing, and to optimise the mass production, the scientific software and the subsystems were embedded in a 'Data Management and Command System'. The main functions of this system were as follows:

(1) to activate the different parts of the reduction software according to a pre-set operation scheme. The organisation of commands allowed the activation of the subsystems sequentially or individually, on one or several data sets. Special care was taken to verify the feasibility of the commands, in accordance with the operation scheme and the available computer resources;

(2) to manage the scientific and operational data within a data bank on disks and magnetic tapes. The Data Management and Command System identified and extracted from the data bank the data necessary to execute a task and stored the results. The storage resources were divided into different areas in order to avoid interferences between parts of the main reduction. A copy of the files was created and managed in a separate archival area;

(3) to allocate dynamically and control the computer resources needed by the subsystems during operation;

(4) to update the consultation area with new results and to manage the access requests. This consultation facility of intermediate and final results through the Data Management and Command System was intensively used by the CERGA team in Grasse, throughout the reductions, for the scientific control of the processing. It was through this facility that all off-line tasks received the intermediate data that was needed;

(5) to allow the debugging, correction of errors, implementation and tests of modifications or software improvements in a separate resource environment.

The size of the Data Management and Command System was about 150 000 lines of code, including comments. It represented roughly half of the complete reduction system, the other half covering the preparation tasks and the scientific software. It managed 4000 high-density magnetic tapes and used 1.5 Gigabytes of disk space. The reduction system worked on the multi-user CDC 992 of the CNES computing centre at Toulouse. The complete reduction of the Hipparcos data at CNES used approximately 1500 hours of CPU time of this 24 Mips machine.

Data Reception and Preparation

The data received from ESOC was first checked for consistency, then divided into different files for subsequent use. From ESOC quality flags the limits of the data sets were automatically determined. After a few months, the test parameters were made less stringent because it appeared that the flagging of possibly bad data was sometimes unnecessarily severe. After this, the quantity of data retained by FAST and NDAC was quite comparable. The preparation of data included the following items:

(1) a full description of the data set was made containing the observation times, the stars observed, the satellite gas jet actuation times and durations of telemetry gaps, eclipse times, the relation between on-board and on-ground times in the form of a third-order polynomial and some statistics. This 'Mission Control File' was used by all FAST software modules throughout the reductions;

(2) several files included all information concerning the stars observed in the data set: in particular *a priori* geometric and apparent positions of stars at their times of observation in the great circle reference frame, various partial derivatives with respect to astrometric and calibration parameters, reference astrometric parameters to be improved, grid coordinates derived from coil-currents, etc. These files were updated at each iteration;

(3) photon counts from the main grid, with details of the observing strategy, and the star mapper photon counts with time indicators. These files were used only in the first treatment;

(4) orbit and on-ground attitude files, and minor planet ephemerides.

First Look

In the FAST Consortium's first-look facility at SRON, Utrecht, data from one orbit per week were received from ESOC within a few days after data acquisition. Using a special version of the FAST data reduction software, these data were analysed within 24 hours of reception; relevant results were distributed to FAST, ESOC and INCA. This special processing served a number of purposes:

(1) a check on the integrity of the data produced by the instrument and processed by ESOC. More particularly, it was verified that the data could be processed without problems by the FAST main reduction software;

(2) a check on the correctness of the reduction software. Especially in the early phases of the data reduction, numerous corrections were proposed. In parallel, results were used in the various comparisons between the consortia processing results;

(3) as a result of the great-circle reductions, geometric calibrations of the main field of view, including the basic angle, were obtained. This provided an excellent method of verifying the stability of the instrument. These results were sent to FAST and ESOC; the latter used the calibration, in particular the grid rotation, to improve the real-time pointing of the instrument;

(4) a number of other quick calibrations were performed: geometry of the star mapper field of view, photometry of the main field, photometry of the star mapper, single-slit response of the star mapper. Results of all calibrations were sent to the FAST main reduction to serve as a first approximation to be improved after the availability of the complete set of measurements. Results of some calibrations were also sent regularly to ESOC; for example, the calibrated modulation factors were sent weekly in order to allow ESOC to monitor the focusing of the instrument;

(5) the level of all star signals were compared with their expected values. In particular, those cases where no signal was measured (which could have been due to an

incorrect position in the Hipparcos Input Catalogue) were signalled to the Input Catalogue Consortium.

Calibrations

The various parameters describing the instrument and the instrumental effects on the data had to be calibrated since the quality of the results of the reduction depended critically on the quality of calibration. So, in FAST, great care and much work was devoted to improve the calibrations so as to use the latest and optimised values at each iteration. Conversely, better intermediate solutions were used to re-run the calibration software in order to further improve the values of the parameters. The main quantities which were calibrated fell into one of the following categories:

(1) image dissector tube data: modulation coefficients M_1, M_2, the phases g_1 and g_2 for single stars, and a first photometric calibration as a function of position on the grid and star colour (see Chapter 5). The calibrations were essentially used in the determination the grid coordinates, in the double and multiple system processing, and in the photometric analysis;

(2) star mapper: the shapes and the position of the star mapper grids with respect to the main grid were calibrated with an accuracy significantly better than the precision of the star mapper observations (see Chapter 6);

(3) optical parameters of the instrument: these included the basic angle and the grid-to-field transformation over the main grid. They were performed as a part of the great-circle reduction (see Chapter 9) although the synthesis and analysis of these results as a function of time were an important part of the calibration task.

Most of these parameters were not stable with time (see Chapter 10) so that the incorporation of the calibrations had to be time dependent. In FAST, the following structure was adopted. The full mission was divided into 33 periods whose duration varied from 6 to 40 days depending on the speed of variation of parameters (mainly the basic angle and the field rotation) also taking into account the refocusing times. Two of these periods concerned sun-pointing situations. For each of these periods, the calibrated quantities were presented, whenever suitable, as analytical functions (polynomials) of position on the grid, colour and magnitude of the star, and time. For the star mapper calibrations, needing lower precision, only 9 periods ranging between 6 and 332 days were adopted. All were periodically updated when better data became available.

Mass Treatment

Within the Data Management and Command System, every data set was treated continuously in a single run from the preparation of the data to the great-circle reduction inclusively, proceeding first through the star mapper data processing, and the determination of the attitude, then the image dissector tube data processing and grid coordinate determination and, finally, the computation of grid coordinates. A number of intermediate results used in other tasks were stored in files: for example the results of the image dissector tube data processing were provided to all the off-line tasks, the attitude to the double star processing and minor planet task, data for various calibrations, etc. Results of the image dissector tube and star mapper photon count treatment and the attitude file were re-used in the iteration mass treatment. The latter therefore included, in addition to a simplified preparation task, a re-evaluation of the attitude and of the grid coordinates and a new reduction on great circles.

Synthesis

Whenever a sufficient number of new or iterated great circle results were obtained, they were sorted star-by-star. The data of some 42 000 stars, selected for their brightness, absence of duplicity, and quality and number of observations, and referred to as 'primary stars' in FAST, were used as an input to the sphere solution software (see Chapter 11). Then, the abscissae of the stars were corrected using the re-determined origins of the great circles and used to compute the astrometric parameters of stars. The latter work was shared by CERGA on the CNES computer, and Astronomisches Rechen-Institut (ARI) in Heidelberg. At this level, the results of the double and multiple star processing were introduced in the equations to determine their astrometric parameters.

The ensemble of the mass production and the corresponding synthesis was called a 'run'. Eight such runs were made during the four years of processing. Four iterations were made during the first 18-month data treatment and three involved more data during the global course of the data reductions. In addition, during the final year, several other syntheses were performed without re-doing the totality of mass processing, but introducing various individual improvements after re-processing some data sets by the mass treatment, or in modifying the status of some stars (primary/secondary or double/single).

Off-Line Tasks

In principle, off-line tasks were performed outside the Data Management and Command System, but the input data were prepared by it. The organisation was as follows:

Minor Planets: The data included the abscissae on great circles as determined by the great-circle reduction task, five parameter modulation coefficients, various data describing the reference great circles, and the orbit of the satellite (position and velocity). The files were prepared in CNES and sent to the Bureau des Longitudes where the corresponding analysis was carried out.

Photometry: This task was essentially based on treated photon counts and modulation coefficients prepared by the Data Management and Command System. The photometry task itself was executed on the CNES computer by the CERGA team. Very good communications existed with the Royal Greenwich Observatory (RGO) team of NDAC, so that comparisons of results were made very frequently and allowed a quick convergence towards comparable results.

Double and Multiple Stars: The complex activities related to the double and multiple star reductions was divided into two separate phases:

(a) relative astrometry of stellar systems, in which the relative positions and intensities of components were determined. Again, the data was prepared within the Data Management and Command System, then sent to the relevant groups. In Italy, it was sent to CSATA (Bari) where further preparation of the data was performed, the results being sent to the Istituto Astrofisica Spaziale in Frascati and the Torino Observatory. The results were compared in Torino and formed the first set of results. A second set was obtained in CERGA on the CNES computer. The FAST ensemble of results was evaluated in CERGA and the best results were sent to Lund for the final merging of double and multiple star data. Multiple star reductions were made independently in Torino and CERGA, then checked, and the results which improved the goodness-of-fit in the computation of the astrometric parameters were retained and sent to Lund for merging;

(b) absolute astrometry of stellar systems, in which the astrometric parameters were determined. This was performed in ARI, Heidelberg. In addition, ARI analysed the residuals of the astrometric parameter reduction, and identified some astrometric double stars and 'variability induced movers'. It also contributed to what are called 'stochastic solutions' for unresolved, probably double or multiple, stars.

4.3. Organisation of the Data Reductions in NDAC

The three main groups of data reduction processes, as described in Section 4.1, corresponded in NDAC with three groups that were each responsible for the development, testing and implementation of their part of the reduction software. The connections between the groups were specified in an interface document, and all data transfer was made by 9-track tape (later also by DAT cartridges) in the agreed standard formats.

The responsibilities were divided as follows:

(1) Royal Greenwich Observatory, UK (RGO): Part A (star mapper processing, attitude reconstruction, image dissector tube data processing), photometric reductions and photometric variability investigations;

(2) Copenhagen University Observatory, Denmark (CUO): Part B (great-circle reductions) and experiments with global solution;

(3) Lund Observatory, Sweden (LO): Part C (sphere solution and astrometric parameter determination) and double star processing.

In all cases the same people who had designed, developed and tested the software also implemented the software, which gave NDAC some advantages of flexibility. All software was implemented on small, semi-dedicated computers (MicroVAX, Sun, and HP/Apollo workstations). The number of people involved at any one time has always been small: at RGO it varied between 3 and 4, at CUO between 1 and 2, and at LO only two people were involved. This reduced greatly the need for meetings, as most problems could easily be solved over the phone or later through electronic mail between the two or three people concerned. A great deal of testing took place before the start of the mission through the use of simulated data produced at the RGO.

The biggest strain on the processing of the data was at the RGO, primarily due to the sheer volume of data arriving and its preparation for the data processing. Approximately 120 to 140 man-days were spent getting the data across from tape to disk in the required format. All data was transferred to optical disk for direct access. A total of 120 disks of 2 Gigabytes each were used to store the complete mission data. This process involved also a first visual inspection of the data quality through an interactive display of the gyro data. Display packages for other data streams were also available and used extensively in the beginning of the mission in order to understand some anomalies.

From the moment the data was on disk, the processing was semi-automatic: programmes were created to produce command files for data processing according to log files of the processing done so far. The command files could keep the processing of many data sets for time intervals of up to several days, the only requirement being the supervision of the occasional change of magnetic tape or disk. The data reductions produced small selections of control graphs, showing the quality of the input data and of the reduced data. A log was kept of every process, which was automatically updated at completion. The standard reductions at this stage were completed by writing the necessary files for the great-circle reductions to tape and sending the data to CUO. During the first 18 months of the mission, updated catalogue data was also sent to CUO.

At CUO the tapes with input for the great-circle reductions were read and the data reorganised on disk, allowing easy access for processing. The reduction of a great-circle set was carried out by performing a chain of processes where the first processes read data from the input files, and the last process wrote the output files. During the processing of a set, several intermediate files were created and used, but these files were all deleted after the output results were accepted. The definition and initialisation of the great-circle set and its reduction were controlled by a number of global parameters which could be changed by editing an option file. Each of the programs used in the reduction of the great-circle sets could be run as stand-alone programs, but the large number of files and tapes involved made it difficult, and indeed undesirable, to manage and monitor the great-circle reduction processes manually. A data management system of programs and directories was therefore set up to control the use of the data files and tapes throughout the mission. Utility programs were also available for plotting results such as the attitude updates or the data in any column of the intermediate files versus record number. From CUO the satellite three-axis attitude for the entire mission was delivered to TDAC, combining results from RGO and CUO.

All the input data for the sphere solution and determination of astrometric parameters, performed at LO, could be contained on a few magnetic tapes and the complete set of files could be stored on disk for the processing. Thus no special data management system was required for this task. However, because there was considerable experimentation with the solution programs and the modelling of various effects, some care was needed to preserve the different program versions and the corresponding data sets. Each new run was given a sequential number (the final sphere solution was number 370), which was used to identify successive catalogue versions as well as the programs by which they were generated. At major milestones of the processing, the whole directory tree containing data and programs was saved on tape cartridges. The solution program also generated files of intermediate abscissa data, used as input to the catalogue merging.

The double star processing at LO used much larger data sets as input, including the signal parameters for the individual field of view crossings from RGO, the complete attitude files from CUO, and a number of photometric and geometric calibration files from RGO, CUO, and the sphere solution at LO. The generation of the 'case history files' from these data streams involved some rather complex juggling of data between disks and tapes (replaced by DAT cartridges in the latter part of the mission). This was however only done a few times throughout the mission, and the subsequent solution of double and multiple stars could then be made one object at a time, using dedicated software for the different object types.

M.A.C. Perryman, J. Kovalevsky, L. Lindegren, F. van Leeuwen