4. Pre-Calibration Data Processing

A set of six software programs, or "facilities", convert the *FIRAS* data from raw telemetered data from the *COBE* satellite to averaged groups of interferograms, or "coadds", that can be Fourier transformed into spectra and calibrated. These facilities are, in order of their use: NFS (iNgest_FIRAS_Stripper); FPP (FIRAS_Pre_Processor); FDQ (FIRAS_Data_Qualify); FEC (FIRAS_Extract_Calibration); FSS (FIRAS_Sort_Sky); and FIL (FIRAS_Interferogram_coaddition_Long_spectra). This section describes these facilities and the delivered *FIRAS* data sets which they produce. A detailed description of this *FIRAS* data processing "pipeline" has been given by Read *et al.* (1992), and is included as Appendix E.

Note that in the following discussion, each "science record" contains one interferogram, or "IFG", together with associated information.

4.1. Ingest

Data from the *COBE* satellite were sent in a telemetry stream to ground stations and received at Goddard Space Flight Center. The *FIRAS* data were separated from other data in the telemetry stream by the *FIRAS* stripper, NFS, which produced ten data sets: one called NFS_HKP, the housekeeping data; one called NFS_ANC, the ancillary data; four called NFS_EMF_xx, the engineering mode timing data; and four called NFS_SDF_xx, the science data, including the IFGs. For the latter two data sets, the "xx" indicates one of the four instrument channels, LH, LL, RH, or RL.

Of these data sets, the data delivery includes the housekeeping, ancillary, and engineering mode timing data. The NFS_SDF data sets are further processed by FPP and FDQ, and so are not part of the delivery.

The housekeeping data, NFS_HKP, consist of time-ordered records each storing the engineering status of the *FIRAS* instrument, including the microprocessor and other status words, temperatures at various locations in the instrument and electronics, the Internal Power Distribution Unit voltages and currents, and other miscellaneous engineering quantities. (The cryogenic temperatures were taken by a set of Germanium Resistance Thermometers, or GRTs). These data are in unconverted telemetry counts. They are stripped from the *COBE* telemetry stream in pairs of major frames, and are not synchronized with the science data. The time of the first major frame is in the header section of the record and the time of the second major frame appears at the end of the record.

The raw ancillary data, NFS_ANC, consist of time-ordered records each containing the pair of telemetry major frames corresponding to the raw housekeeping data. For each major frame, the record contains 128 minor frames of status bits packed into one byte per minor frame. The status bits are the MTM scan direction bit, the external calibrator status bit, the MTM scan length, the MTM scan speed, and four microprocessor data ready bits.

The engineering mode timing data, NFS_EMF, consist of time-ordered records of Mirror Transport Mechanism timing information. These data provide a direct measurement of the time between MTM samples. For each channel, the flight microprocessors output the elapsed time between samples for one MTM sweep into a 512-point, 8-bit array. The time elapsed between sample pulses is computed in microseconds as follows:

$$time \;=\; 45 \,+\, (10.0 \cdot counts) \tag{1}$$

where *counts* is the timing data. Engineering mode timing data were taken early in the mission during the post-launch spacecraft checkout and late in the mission during post-cryogen depletion engineering tests. Analyses of these data yielded the MTM sampling rates discussed in Section 2. Fourier analysis of the data yielded the frequencies of the coherent vibrations discussed in Sections 5 and 7.

The NFS_HKP, NFS_ANC, and NFS_EMF data sets are released in their native VAX binary file formats, the record structures for which are given in Appendix H. The filename extensions of these files (*e.g.*, ED_893272157) incorporate the timetag of the earliest data record contained in the files, using the format YY-DDD-HH-MM.

4.2. Quality Checking

FPP and FDQ are the next two facilities in the *FIRAS* data processing pipeline.

The FPP facility has two main functions: to calculate the midpoint of each IFG's collection time, which is used to determine the appropriate spacecraft attitude solution for the IFG; and to perform checks on the quality of the data. To calculate the midpoint of collect time, FPP uses the data transmit time in the science record to extract the mirror transport mechanism (MTM) scan speed and stroke length from the corresponding engineering and status records. The on-board microprocessors for each detector channel telemeter 26 header words of information with every IFG collected. Several of these header words contain counters for the start of the IFG collection and the start of IFG transmission to the telemetry stream. FPP extracts the MTM synchronization flag and the number of sweeps that are averaged together to produce the IFG from the header. It uses these MTM data, the counters, and the actual transmit time for the IFG record to compute the midpoint of the collection time. If there are any anomalies in the values used in the time computation, the IFG science record is flagged as being of bad quality. The IFG is also flagged if there are telemetry data gaps or bad telemetry quality during its collection. Finally, FPP checks the values of the MTM scan speed, MTM stroke length, and the detector gain over the IFG collection time. If these values are not constant, the IFG is flagged. Any record flagged by FPP is not used in further pipeline processing.

Table 4.1 gives the number of science records collected by the *FIRAS* instrument in each of the four channels over the entire 10 month *COBE* mission, and gives the number failed by FPP. The data are divided into three groups according to the position of the external calibrator (XCAL). Calibration data records have the XCAL in the sky horn, while sky data records occur when the XCAL is in its stowed position and the *FIRAS* instrument is looking at the sky. Some data were collected while the XCAL was in transit between the sky horn and its stowed position; these records are not used. Also, a special data mode called "fake-it" was used for about 10% of the records to conduct a variety of engineering tests; these records are also not used in further processing.

Table 4.1: *FIRAS* Facility FPP

Channel	LH	LL	RH	RL	All
Total Science Records	590937	590926	591005	587637	2360505
Records Failed by FPP	22118	22865	23315	22684	90982
Fake-it Records	59637	58797	59136	58348	235918
XCAL Transit	430	430	312	291	1463
Records Eliminated Before FDQ	82185	82092	82763	81323	328363
Cal Records Passed by FPP	61769	61755	61915	61860	247299
Sky Records Passed by FPP	446983	447079	446327	444454	1784843

The FDQ facility has three main tasks: to continue the data quality checking begun by FPP; to associate the science records with appropriate engineering data (such as temperatures, voltages, and currents); and to determine the spacecraft attitude solution (and thus, a pixel number) for each IFG.

FDQ checks the microprocessor, engineering, status, and attitude quantities associated with each IFG, and sets quality flags for the records appropriately. Most of the data quality checks are for engineering quantities, but very few records failed any of these during the mission. The principal failures are for the microprocessor status word, saturated sample counts, elevated glitch rates, and attitude.

The microprocessor status word contains 16 bit flags which indicate errors or anomalies detected by on-board programs during the collection of IFGs. These errors include deglitcher math overflow, command overrun, data buffer clear error, sync and ADC pulse time-outs, illegal instruction fetch, sync errors, ADC buffer over-run, sample add and divide overflows, group coadd overflow, and group divide overflow. If any of these bits are set, FDQ gives the IFG a bad quality flag. The saturated sample count is defined as the number of ADC samples greater than or equal to 95% of the maximum ADC range. If this count is too high, the IFG is also failed. Data collection by the *FIRAS* instrument is affected by cosmic ray hits on the detectors, or "glitches". A glitch rate is computed for each IFG from a telemetry data word containing the glitch count for the IFG. A glitch rate that is too high causes FDQ to fail the IFG.

FDQ also assigns bad data quality to science records based on attitude quantities. If no attitude solution is available for a particular IFG, that record is failed. It is also failed if the attitude is such that the IFG might have been effected by emission from the Sun, the Moon, or the Earth limb. Because the attitude had no demonstrable effect on data taken when the XCAL was in the sky horn, calibration IFGs are not checked for attitude failures.

Table 4.2 gives the number of calibration and sky data records failed by FDQ for each of these reasons.

Table 4.2: *FIRAS* Facility FDQ

Channel	LH	LL	RH	RL	All
Saturated Sample Count	4581	5499	4462	4938	19480
Microprocessor Status Word	685	317	585	307	1894
Glitch Rate	627	181	444	271	1523
Cal Records Failed by FDQ	5053	5903	4905	5339	21200
Cal Records Passed by FDQ	56716	55852	57010	56521	226099
Saturated Sample Count	26694	44098	26206	44242	141240
Microprocessor Status Word	6400	6672	6893	6421	26386
Glitch Rate	3619	3731	3541	3779	14670
Sun, Moon, or Earth Limb	7814	7807	7796	7788	31205
No Attitude Solution	569	569	566	329	2033
Sky Records Failed by FDQ	41281	59640	41206	58353	200480
Sky Records Passed by FDQ	405702	387439	405121	386101	1584363

FDQ produces five data sets: FDQ_SDF_xx, which contain the IFGs and the associated attitude information, and where "xx" is one of the instrument channels LH, LL, RH, or RL; and FDQ_ENG, which contains the engineering data.

The time-ordered FDQ_SDF records are created from the IFGs collected by the on-board microprocessors and from spacecraft attitude information corresponding to the midpoint of collect time of the IFGs. The engineering data, FDQ_ENG, consist of time-ordered records each storing the engineering status of the *FIRAS* instrument at a particular time. The status includes the microprocessor status words, temperatures, voltages, and currents. Both FDQ_SDF and FDQ_ENG records are stored in files each containing one day of *FIRAS* data.

Because the science and engineering data are collected asynchronously in order to avoid coherent crosstalk, FDQ must properly link them together. To do so, the science IFG records are grouped according to the proximity of the midpoint collection times. Each such group may have from one to four science records, up to one from each of the four channels. The two housekeeping major frames whose times bracket the average time of the science record group are found. FDQ converts the data in these housekeeping records from microprocessor counts to physical units, and then interpolates them to determine values for each engineering quantity at the average science record group time. This time is assigned to the new FDQ_ENG engineering record, and is also put into each of the FDQ_SDF science records in the group. The science record times are put into the engineering record, so that the science and engineering data are doubly cross-indexed.

The FDQ_SDF_xx and FDQ_ENG data sets are released in their native VAX binary file formats, the record structures for which are given in Appendix H. As with the NFS data, the filename extensions of these files (*e.g.*, ED_893280000) use the format YY-DDD-HH-MM; however, because the FDQ files each contain one day's worth of data, the hour and minute are always zero.

4.3. Sorting

At this stage of the *FIRAS* pipeline, the calibration data and sky data are separated and processed independently. The FEC (for calibration data) and FSS (for sky data) facilities perform two functions: to sort the individual IFGs into ensembles for coaddition; and to provide further checks for data that are unsuitable for further use.

During the *FIRAS* mission, there were 16 intervals of typically two to three days duration when calibration data were taken (Table 3.1). During calibrations, series of time-ordered IFGs were obtained; these series were separated either by: a change in the commanded temperature for at least one of the four temperature-controllable bodies (the external calibrator (XCAL), the internal calibrator (ICAL), the sky horn, or the reference horn); or by a change in the commanded detector bias voltage. Each such series constitutes a coaddable group of calibration IFG records. The averages of the temperatures of the four controllable bodies are computed for each series. An IFG within the series is removed from the ensemble if the temperatures of any of the four controllable bodies associated with that IFG differ from the average for the series by more than a specified tolerance. IFGs with deviant temperatures are classified as unstable and are removed from further pipeline processing. Many different sets of specified tolerances were used for the calibration data; as an example, the set used for the time range 900780133 to 900782047 was: XCAL .1%; ICAL .05%; sky horn .2%; and reference horn .05%.

Most of the calibration records which fail stability tests are at the beginning of a series, because a controllable body which has just experienced a commanded temperature change requires a finite time to relax to its new state. The other major cause of stability failure occurred following unplanned transient heatings of the sky horn by the bright earth limb near the end of the eclipse season, when horn temperatures were 2.7 K. The calibration data were screened to avoid these events. Some records were unstable because of transients which occurred during the initial five minutes or so of a series defined by a change in commanded bolometer bias voltage.

Each coaddable group of IFGs which pass the stability tests is subdivided if there is a change in MTM scan mode and further subdivided if there is a change in science gain. Groups containing more than 100 IFGs are divided into smaller groups of approximately equal size. The number of segments is chosen to maximize the size of the subgroups. For example, a coadd group containing 700 IFGs would be divided into seven time-ordered groups of 100 IFGs each, a coadd group containing 380 IFGs would be divided into four groups containing 95 IFGs each, etc.

Table 4.3 gives the number of calibration IFGs failed and passed by FEC, as well as the number of coadd groups the facility created.

Table 4.3: *FIRAS* Facility FEC

Channel	LH	LL	RH	RL	All
Cal Records Failed by FEC	19275	18918	19459	19187	76839
Cal Records Passed by FEC	37441	36934	37551	37334	149260
Cal Coadd Groups Passed by FEC	771	762	767	768	3068

To allow the FSS facility to form coaddable groups of sky IFGs, the FDQ data are separated into 11 mission time periods according to various events and changes in the instrument state. The time ranges and descriptions of the mission periods are given in Table 4.4.

Table 4.4: *FIRAS* Mission Periods

Start Time	Stop Time	Description of Interval
89-326-1130	89-327-2359	First light
89-328-0000	89-343-0151	First ICAL nulling
89-343-0152	90-019-0204	MTM uses position mode through SAA
90-019-0205	90-080-0114	Horns commanded from 2.70 K to 2.75 K
90-080-0115	90-128-2359	MTM uses power off through SAA
90-129-0000	90-139-1534	Eclipse season starts
90-139-1535	90-193-1849	Horns commanded to 6 K
90-193-1850	90-207-1103	Horns commanded to 4 K
90-207-1104	90-208-1119	Sky horn calibration, XCAL out
90-208-1120	90-220-0459	Horns commanded to final temperature
90-220-0500	90-264-0936	XCAL placed under temperature control

FSS is run on each mission period in turn, for all four channels. It first checks the sky IFG records for the following conditions, rejecting those which do not qualify: a solar aspect angle greater than 91.2° ; an earth limb aspect angle greater than 87° ; a lunar aspect angle greater than 22° ; an ICAL temperature within 2 mK of one of several specified temperatures; a dihedral temperature less than or equal to 5.5 K; and a science mode of four (the standard operating science mode).

Table 4.5 gives the number of sky IFGs passed and failed by FSS, as well as the number of coadd groups it creates for each of the four channels.

FSS then sorts the acceptable sky IFG records into coaddable groups. It first sorts by pixel number (location on the sky), then by MTM scan mode (one of SS, SF, LS, or LF), then by ICAL temperature (within 2 mK of the specified ICAL temperatures, thus forming "bins"), and, finally, by several ranges of dihedral temperatures. Coadd groups of more than 100 IFGs are subdivided into smaller sized subgroups in the same way as calibration coadds are subdivided. Table 4.6 gives the ICAL temperatures (forming the center of the "bins") and the dihedral temperature ranges for each of the 11 mission periods.

Once the coadd groups have been formed, FSS forms additional groups of "neighbor" IFGs.

Table 4.5: *FIRAS* Facility FSS

Channel	LH	LL	RH	RL	All
Earth Limb Angle <87.0	10914	10763	10939	10706	43322
Wrong ICAL Temperature	7415	7099	7370	5250	27134
Sun Angle <91.2	6745	6764	6723	6768	27000
Wrong Science Mode	718	717	744	827	3006
Dihedral Temperature >5.5	381	380	371	332	1464
Sky Records Failed by FSS	23745	22586	23697	20752	90780
Sky Records Passed by FSS	381957	364853	381424	365349	1493583
Sky Coadd Groups Passed by FSS	78577	77562	78547	77592	312278

Table 4.6: *FIRAS* Facility FSS Temperatures

Start Time	Stop Time	ICAL Temperatures	Dihedral Temperatures
89-326-1130	89-327-2359	2.789	2.14,2.5,3.1,3.7,4.3,4.9,5.5
89-328-0000	89-343-0151	2.758,2.763,2.789	2.02,2.5,3.1,3.7,4.3,4.9,5.5
89-343-0152	90-019-0204	2.759,2.771	2.14,2.5,3.1,3.7,4.3,4.9,5.5
90-019-0205	90-080-0114	2.758,2.771	2.14,2.5,3.1,3.7,4.3,4.9,5.5
90-080-0115	90-128-2359	2.758,2.771	1.98,2.5,3.1,3.7,4.3,4.9,5.5
90-129-0000	90-139-1534	2.758,2.770	2.0,2.5,3.1,3.7,4.3,4.9,5.5
90-139-1535	90-193-1849	2.7455,2.755,2.768	2.03,2.5,3.1,3.7,4.3,4.9,5.5
90-193-1850	90-207-1103	2.746,2.757,2.769	2.01,2.5,3.1,3.7,4.3,4.9,5.5
90-207-1104	90-208-1119	2.757,2.769	2.01,2.5,3.1,3.7,4.3,4.9,5.5
90-208-1120	90-220-0459	2.758,2.770	2.0,2.5,3.1,3.7,4.3,4.9,5.5
90-220-0500	90-264-0936	2.758,2.771	2.0,2.5,3.1,3.7,4.3,4.9,5.5

For each coadd group, FSS finds sky IFG records with the same MTM scan mode, ICAL temperature within the same bin, and dihedral temperature within the same range from the eight surrounding, or "neighboring" pixels. It then computes the mean Galactic latitude of the original coadd group, and orders the neighbor IFGs by increasing absolute values of the difference between the Galactic latitude of each neighbor record and this mean.

Unlike calibration data, IFGs in sky coadd groups are not tested for XCAL, sky horn, or reference horn temperature stability. Analysis has shown that variations in the XCAL temperature do not have a detectable effect on the calibrated spectra. Because the Mission Periods are defined so that the horn temperatures do not vary greatly within them, no sky coadd contains records with greatly different horn temperatures. Horn temperature differences between coadds with the same sky pixel number obtained in different Mission Periods are accounted for by the calibration model (Section 5). Residual post-calibration effects are corrected by destriping (Section 6).

The FEC_SSCAL_xx and FSS_SSSKY_xx records, which are the FEC and FSS products and where "xx" denotes one of the channels LH, LL, RH, or RL, are released in their native VAX binary file formats, the record structures for which are given in Appendix H. The filename extensions of these files (*e.g.*, ED_8935312_8935323) incorporate the timetags of the earliest and latest data records contained in the files, using the format YY-DDD-HH. These files do not contain any science or engineering data. Instead, they are index files that refer back to the FDQ_SDF and FDQ_ENG files.

4.4. Coaddition

The FIL facility coadds the IFGs by performing the following functions: reading the data; checking the data for quality and instrument state consistency within each coadd group; creating and subtracting primary and, possibly, secondary templates; removing the effects of cosmic ray hits on the detectors (called "deglitching"); checking the shape of the IFGs for consistency within each group; coadding the science and engineering data using a weighting scheme based on glitch rates; creating, as required, "FS" or "FL" scan mode IFGs from SF or LF scan mode IFGs; and subtracting a baseline from each coadded IFG.

FIL first uses the FEC or FSS products to retrieve the FDQ_SDF science records (including the IFGs) and the FDQ_ENG engineering records. Members of the original coadd group (excluding neighbors) are always read. If FIL is operating on sky (FSS) data and if there are fewer than 12 IFGs in the original group, then neighbor IFGs are read (if they exist) until a group of 12 is formed.

The neighbor IFGs are not coadded together with the "original" IFGs. Instead, they are used for the template formation and deglitching process. Analysis shows that a group of at least three IFGs is needed to form a reliable template. Since many sky data coadd groups contain only one or two IFGs, the use of neighbors permits the IFGs in these groups to be included in template formation and deglitching, and thus prevents their loss. (The coadd groups of calibration data created by FEC always contain five or more IFGs).

FIL next checks that none of the IFGs in a coadd group were marked as bad by a previous pipeline facility, and that all of them are in the same instrument mode. Verification of the instrument mode includes checking that attitude quantities and operating conditions such as channel, scan mode, science mode, and temperature controller gain status bits are the same for all IFGs. For engineering quantities such as bolometer voltage, temperatures of the four controllable bodies, and bolometer temperature, FIL calculates midaverages of the coadd group. IFGs with measures not within tolerance ranges of these midaverages are discarded; these tolerances are defined in the reference data set FEX_CTH.TXT (Section 9.1). Specifically, bolometer temperature tolerances are set at the $\frac{1}{2}$% level so that the dispersion in the bias voltage (which is highly correlated to the bolometer temperature) will be less than 1%, while GRT tolerances are chosen to give calibrated flux changes of less than 2% in the worst case.

If in the course of performing these checks the number of remaining good IFGs (including neighbors) falls to less than three, then the one or two remaining IFGs are considered to be "too few to coadd", and are marked as bad.

Following the quality and instrument consistency checks, if the number of remaining good IFGs including neighbors is less than or equal to eight, then all of them are used in subsequent processing. If the number of good IFGs excluding neighbors is greater than or equal to eight, then no neighbor IFGs are needed and all of them are marked as bad. If, on the other hand, the number of good IFGs including neighbors is greater than eight, then enough neighbor IFGs are marked as bad to bring the total down to eight. The question of which neighbors to mark bad is decided as follows. FIL calculates the mean Galactic latitude of the good IFGs excluding neighbors. If this mean is less than 10.0, then the neighbor IFGs with the greatest absolute values of the difference of their Galactic latitudes and the mean are marked bad. If the mean is greater than or equal to 10.0, then the neighbor IFGs with the highest glitch rates are marked bad.

In order to normalize data accumulated under differing operating conditions, each IFG is next divided by the real-valued commanded preamplifier gain and the number of onboard sweeps of the MTM, since the digitized voltages are proportional to these quantities. FIL also removes the "dither", a randomly generated value, different for each IFG, which was added to each point of the IFG during its collection. It does so by finding the median value of each IFG and subtracting it.

A robust mean is then constructed as a template by calculating the pointwise midaverage of the coadd group. This mean is the average of the middle two quartiles of the coadd group values at each point. This so-called "primary template" is subtracted from each IFG.

To remove variable signals from gas and dust in the Galactic plane, a "secondary template" is then calculated from the coadd group of interferograms. The IFGs are temporarily aligned so that the signal at the zero path difference is positive in each case; after this is done the secondary template is calculated in the same way as was the primary template.

To ensure that secondary templates are removed only when necessary, secondary template subtraction occurs only if the amplitudes and signal-to-noise ratios of both the primary and secondary templates exceed reference thresholds as specified in the reference data set FEX_CTH.TXT. In that case, suitable multiples of the secondary template are subtracted from each IFG so that the total of the absolute values of the differences between the individually scaled secondary template and the IFGs is at a minimum for 21 points centered at the zero path difference. These secondary template subtraction cutoffs were deemed necessary because this process was seen to influence performance of the deglitching algorithm (see below); the deglitcher removed too few or too many supposed glitches near the peak of the IFG in some cases if the secondary template was subtracted.

A transient response from the onboard digital filters is next subtracted from the first 128 points of each IFG using a least squares fit. The transient is very small after the first few points.

Next, a pattern recognition algorithm is applied to remove contamination due to glitches on the detectors, which are the dominant noise source. For each IFG, the deglitching algorithm uses a robust noise estimator which is the maximum of 1.25 times the median absolute value of the IFG and the threshold bit noise. In an iterative process, the largest value of the 512 points of the template-subtracted IFG is found. If the absolute value of the ratio of this value to the IFG noise is greater than 3.7, a glitch is deemed to exist at that point; a three-point parabolic fit is used to determine the "true" peak position of the glitch, and the glitch profile whose peak is closest to the glitch peak is subtracted. (See Section 9.2 for a description of the glitch profiles). The height of the profile subtracted is scaled by .2 if the ratio above is greater than or equal to 5.5, and by .7 if the ratio is less than 5.5 but greater than or equal to 3.7. The process repeats until all glitches are removed (that is, no point has a ratio greater than 3.7). This technique is based on the *CLEAN* algorithm and has been described in detail by Isaacman, Read, and Barnes (1992), which is included in this document as Appendix F.

Following deglitching, "shape consistency checking" is used to reject anomalous IFGs whose overall shape or noise properties differ significantly from those of the coadd group as a whole. The noise (sigma) of each IFG is defined to be 1.25 times the median of the absolute values of the deviations from its median, and the median sigma of the coadd group is defined to be the larger of the median of these numbers and a value denoting a deviation of one bit. The ratio of each IFG's sigma to that of the ensemble is calculated, and IFGs with ratios above or below values of 1.5 and 0.5 respectively are considered to have "high noise" or "low noise", and are marked as bad. For each remaining IFG, the ratio of its value at each point to the median sigma is then calculated; it is marked as bad if more than six points have a ratio greater than six.

Table 4.7 gives the numbers of calibration and sky IFG records and coadd groups failed for various reasons, as well as the number of IFG records and coadd groups (which became coadded IFGs) that passed.

FIL now coadds the science and engineering data using a weighting based on the glitch rates of the IFGs being coadded. Even after being deglitched, the high glitch rate data have higher noise than the low glitch rate data. This is not surprising since the glitches are not perfectly determined, the glitch profiles are not perfect, and there are a large number of glitches that are too small to be detected by the deglitching algorithm but which still contribute to the overall noise.

The glitches are essentially delta functions of energy incident on the bolometers, and so add power at all frequencies. Because of the filtering of the data and the response of the bolometers, more of the glitch noise shows up at low frequencies than at high frequencies. The arrival times of the cosmic rays are not correlated with the phase of the *FIRAS* scans, and so do not bias the data in any way. The glitches add only noise. Furthermore, because the *FIRAS* pipeline handles calibration data the same way that it handles sky data, any bias introduced to the data by the deglitching algorithm occurs in both the sky and calibration data, and consequently is removed during calibration.

Since the high glitch rate data have a higher noise, these data are deweighted with respect to the low glitch rate data. Once the coadded IFGs were converted to spectra and calibrated, an average variance per frequency sample for each spectrum could be computed:

$$\langle var(i) \rangle_\nu \;=\; \frac{nifgs(i)}{nfreq} \sum_{\nu=0}^{nfreq} \left(\frac{variance(\nu, i)}{D(\nu)^2} \right) \qquad (2)$$

where $nifgs(i)$ is the number of IFGs in spectrum i and $nfreq$ is the number of frequency points in the spectrum. The D is from a previous calibration (Section 7.1.1). These

Table 4.7: Records failed and passed in the *FIRAS* Facility FIL.

Channel	LH	LL	RH	RL	All
Bolometer Temperature	2383	2864	2004	2294	9545
Bolometer Voltage	914	477	528	295	2214
Noise Too High	151	468	468	612	1699
Mismatched Science Mode	135	130	147	149	561
All Other Failure Reasons	26	139	160	177	502
Cal Records Failed by FIL	3609	4078	3307	3527	14521
Cal Coadds Failed by FIL	12	11	9	8	40
Cal Records Coadded by FIL	33832	32856	34244	33807	134739
Cal Coadds Produced by FIL	759	751	758	760	3028
Bolometer Temperature	33943	40474	36230	39722	150369
Sky Horn Temperature	12205	12856	12488	12851	50400
Bolometer Voltage	17463	8231	10769	4643	41106
Reference Horn Temperature	1887	2001	1990	1944	7822
Too Few to Coadd	6120	6274	6122	5893	24409
Temp Cont Gain Status Bits	3120	3002	3244	3149	12515
Noise Too High	1608	3127	3873	5479	14087
All Other Failure Reasons	347	412	867	629	2255
Sky Records Failed by FIL	76693	76377	75583	74310	302963
Sky Coadds Failed by FIL	8833	9223	9034	8658	35748
Sky Records Coadded by FIL	305264	288476	305841	291039	1190620
Sky Coadds Produced by FIL	69744	68339	69513	68934	276530

average variances were then modeled as a linear function of the glitch rate:

$$\langle var(i) \rangle_\nu \;=\; slope \cdot glitch\_rate(i) \,+\, intercept \tag{3}$$

The results of this modeling are given in Table 4.8.

Table 4.8: Dependence of Variance on Glitch Rate

Scan Mode	Slope	Error	Intercept	Error
LHSS	1.5191	0.0954	0.8917	0.0075
LHSF	0.7267	0.1775	0.9526	0.0137
LHLF	0.2083	0.0686	0.9840	0.0054
LLSS	0.9034	0.0166	0.6037	0.0080
LLSF	1.1911	0.0595	0.6090	0.0207
LLLF	0.7500	0.0194	0.6825	0.0088
RHSS	0.3181	0.0151	0.8078	0.0096
RHSF	0.2141	0.0304	0.8748	0.0190
RHLF	0.0967	0.0123	0.9389	0.0081
RLSS	1.4353	0.0259	0.4982	0.0096
RLSF	0.8577	0.0493	0.7115	0.0181
RLLF	0.5659	0.0246	0.8027	0.0088

The slope and intercept from this fit was used to renormalize the number of IFGs in the spectra:

$$weight(i) \;=\; \frac{nifgs(i)}{slope \cdot glitch\_rate(i) \,+\, intercept} \tag{4}$$

The net effect of this renormalization is to redistribute the weight within and to decrease the total variance in the data set. This occurs because the decrease in variance of the low glitch rate data more than offsets the increase in variance of the high glitch rate data. The total decrease in variance is ~2%.

The good IFGs in the coadd group are coadded by: determining the weight for each IFG; calculating the sum over the coadd group of each IFG multiplied by its weight; and dividing by the sum of the weights. The primary template and, if appropriate, a multiple of the secondary template formed by averaging the individual factors used in secondary template subtraction are added back to this result to form the “coadd”.

The engineering and attitude data associated with the IFGs are coadded in the same manner as the IFG itself. Each coadd is thus accompanied by mean instrument voltages, temperatures, etc., and a mean sky position. Note that this quantity need not be at the nominal pixel center.

Because there was relatively little low frequency data in the SF scan mode taken during the mission, it proved impossible to create a calibration model for it alone (Section 5). A scheme was devised to combine "decimated" calibration coadds for low frequency SF data with "truncated" ones for low frequency LF data to permit the computation of a calibration model that could be applied to both. At this point in its processing, FIL creates both calibration and sky coadds for these two new modes, which are given the designations "FS" for decimated SF and "FL" for truncated LF. The truncated FL coadds are simply the first 128 points of the 512 point LF coadds. The decimated FS coadds are generated using the following algorithm: discard the first three points of the SF coadd; for points 1 to 127 of the FS coadd, take the next four points of the SF, add them together, and divide by four; point 128 of the FS is set to point 512 of the SF.

To remove the effects of internal defocussing of the instrument at the extrema of the MTM sweep, a fourth-order polynomial baseline is fitted to the coadd using a least squares fit and subtracted (see Section 9.2 for details of the baseline function). Last, the transient response from the onboard digital filters is subtracted from the first 128 points of the coadd using a least squares fit (as was previously done for each of the individual IFGs.)

Because so little LS scan mode data was taken during the mission, no calibration model exists for it. Thus, of the original 16 possible channel/scan mode combinations and of the four new low channel/FS or FL scan modes, FIL creates data sets for LHSS, LHSF, LHLF, LLSS, LLFS, LLFL, LLLF, RHSS, RHSF, RHLF, RLSS, RLFS, RLFL, and RLLF. The project data release includes the time-ordered coadded calibration interferograms and the pixel-ordered coadded sky interferograms for each of these channels and scan modes.

The FITS binary and extension headers for these data products are given in Appendix G; they include the names and descriptions of each of the data fields of these data sets.