Report on an LWS01 Processing Problem

T W Grundy, UKIDC, RAL

6th May 2004

Introduction

The LWS OLP 10 Scientific Validation Report describes the checks made on the LWS OLP 10 test case data products to determine whether they are scientifically valid. The test cases used a sample of products from the L01, L02 and L04 observing modes. L03 was not examined and the L03 AAR products are not considered scientifically valid. The three modes tested were found to be scientifically valid.

In 2003, the validity of the L01 mode was re-examined, for a subset of the observations not previously investigated as part of the validation process, following the discovery of a problem with the detector responsivity drift-correction. The problem was found to be related to a bug in the processing software, which is introduced at SPD level. Subsequently, the bug then caused a malfunction in AA level processing, resulting in a degradation of the scientific quality of the data in the products affected.

An investigation was made into this problem (referred to as the 'scan-numbering problem'), to determine the cause and a solution. This is described in detail in the sections below. The procedure devised to find and correct the products affected is subsequently explained. Finally, the procedure used to test the corrected products is outlined. The reprocessed products are shown to meet the standards determined in the OLP 10 scientific validation report.

Introduction to the Scan-Numbering Problem

One focus of the work at the UKIDC in 2003/4 has been on converting LWS calibration data products into standard pipeline-style products. This is referred to as COIF conversion, and involves extracting the contents of the raw data files (LXER) into standard ERD files. These can then be processed using the standard OLP pipeline software, producing SPD and AAR products. The L01 mode AAL processing was found to fail in certain cases, so a more detailed investigation was made to determine the cause of the problem. During this investigation, it was noticed that similar problems can also occur with some ordinary L01 observations. However, no L01 observations were found to cause unsuccessful terminations of the processing software, so there hasn't previously been an obvious clue that something was wrong. Additionally, the final AAR products obtained from both COIF and L01 data frequently appear normal, so no problems have ever been reported during quality control checks. In a minority of cases, it has been found that the processing problems cause a serious impact on the final science data, such that the resulting LSAN files become virtually useless. The total number of ordinary L01 observations known to be affected in OLP10.1 is 277.

The next two sections describe the details of the problem, and the correction that was devised.

Details of the Problem in the Pipeline

The LWS pipeline is split into two stages: SPL and AAL. SPL is the standard processing pipeline, and deals with the low level details of converting edited raw data (ERD) into basic science information. There are many stages, most of which are explained in detail in the LWS handbook. One stage that is not mentioned in the handbook is the assignment of scan numbering. This is part of the responsivity drift correction algorithm, which is unusually split between the two pipelines. In the SPL, the movement of the grating or Fabry-Perot is analysed, and scan numbers are assigned to each full sweep of the optics mechanism in use. For L01 products, this is the grating. The grating commanded and measured positions were recorded on board LWS, and included in the ERD. A standard observation consisted of alternate ascending and descending passes or 'scans' of the grating, made in order to cover the wavelength range required in the observation. The first scan is always 'forward', which corresponds to an increase in the grating position readout. A new scan begins every time the grating changes direction. In the SPL, the scan numbering is simply assigned for convenience – it is not used by any of the scientific algorithms implemented to produce the SPD products.

The AAL is the auto-analysis pipeline, and is responsible for the calibration of the basic science data in the SPD products. The first stage of the pipeline is the analysis of the illuminator flashes. This uses only the LIPD illuminator data, and produces the LIAC illuminator summary file, which is hence unaffected by any problem in the LSPD file. The second stage is the rest of the responsivity drift correction algorithm, which was started in the SPL with the assignment of scan numbering. This is itself split into three parts.

1. Each scan is averaged, after removal of invalid data points (according to the SPD status words), to produce an 'average signal per point' value.

2. In each calibration group (the period between each illuminator flash or raster point change), the set of 'average signal per point' values is analysed to look for a time-dependent drift trend. This involves fitting a first-order polynomial to the values and times for each scan. A pair of correction coefficients is hence determined, to be applied to the group.
3. The correction is applied to the photocurrent data, using the calculated coefficients.

The problem is that the first stage depends on the correct assignment of scan numbers in the LSPD file. In a normal case, the scan number changes at the end of each grating scan, and the algorithm works as it was intended. However, if the scan number changes in the middle of a grating scan, then the average value calculated is not suitable for comparison with that for a normal, full scan. The following example shows what can happen.

TDT 84900612 was one of the worst affected observations. It consists of 3 full grating scans (2 forward and 1 reverse), and 1 short grating scan (reverse). These should be numbered from 0 to 3; however the LSPD file contains scans numbered from 0 to 4. The first scan is covered by two scan numbers (0 and 1), and the change from 0 to 1 happened very early in the scan. To see how this affects the drift correction, any one detector can be analysed. The following plot shows the average signal per point values, calculated for each of the incorrectly numbered scans.

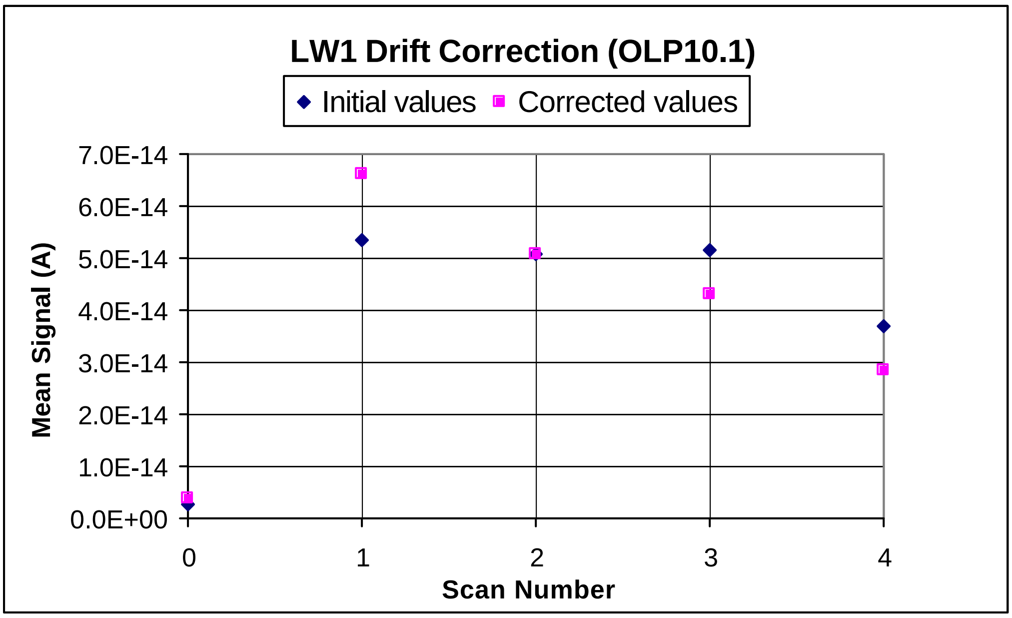

The scans numbered 1, 2 and 3 represent the bulk of the data, and are almost 3 full scans. The first few data points are in the value for scan 0, and the short scan is number 4. You can see that the values (dark blue) for scans 1 to 3 are very similar, indicating that the detector responsivity drifted only a small amount over the course of the observation. However, due to the bad scan numbering, there is an invalid data point for scan 0. Also, the fact that it has occurred at the start of the observation means that scan 4 is no longer considered by the pipeline to be a short scan, as it is comparable in length to scan 0. Therefore, the second stage of the algorithm uses all 5 data points to perform the fit. The resulting drift slopes have then been applied to the data points to produce the corrected data points (magenta) in the graph. You can see that the values for scans 1 to 3 are no longer similar, and the values now imply that the detector has drifted downwards in responsivity. The result is that, contrary to the intention of the drift correction, the scans have been pushed further apart in responsivity.

The following plot shows the un-averaged LSAN data for TDT 84900612.

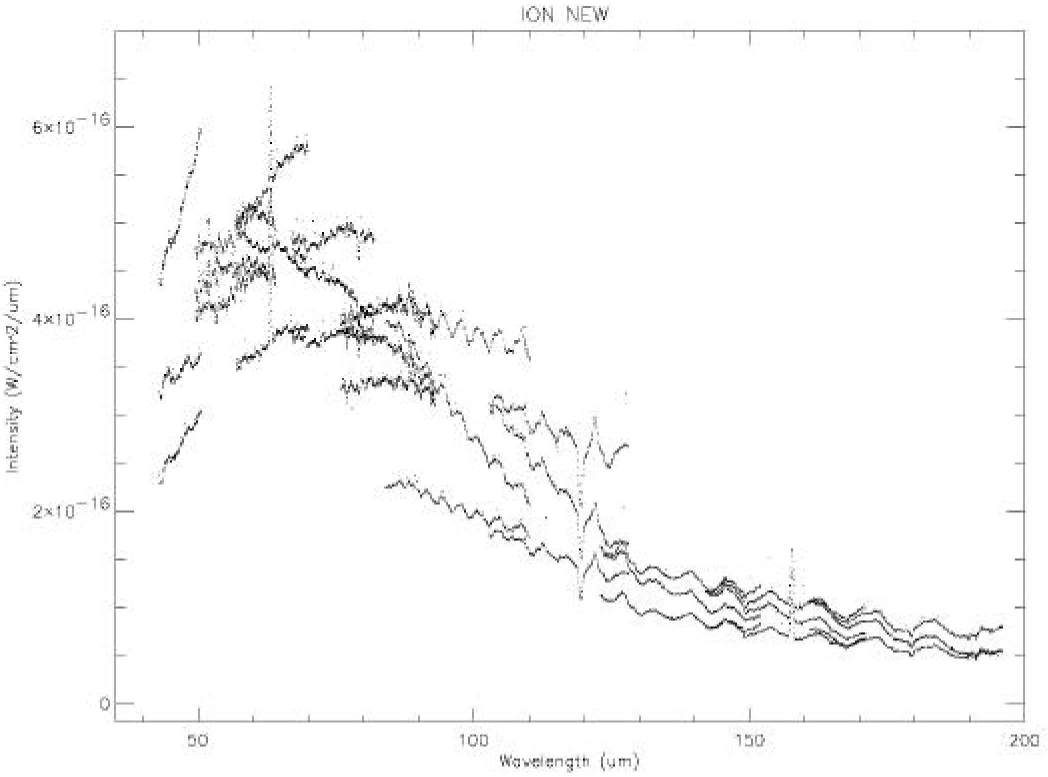

The scans show poor agreement with each other in all detectors, although some are worse than others. LW1 covers the range from 84 to 110 microns in the centre of the plot, and is especially badly affected.

Fixing the SPD

The exact cause of the scan-numbering problem is not known. The most likely cause is that occasionally, small sections of raw data are missing from the ERD products, due to telemetry drops or other transmission / reception problems. It may be that a large enough section of missing data could cause the scan-numbering algorithm to malfunction, and start a new scan prematurely. The problem has only ever been found in the first grating scan in any given raster point (scan numbers always start from 0 in each raster), so it seems that the algorithm is robust once a full scan has been completed, and a change of scan direction has been observed.

As the scan numbering doesn't affect any other element of the SPL, the simplest solution is to fix the LSPD product. Before the discovery of the problem in LWS01 observations, the same problem was found in LWS99 COIFs. Some of these observations were made in a very similar way to LWS01s, and an algorithm already existed to fix these products. An attempt was made to apply the same fix to a normal LWS01, and was found to be successful. The fixing algorithm is based on different assumptions to those in the SPL, and will ignore any sections of missing data. The operation is simple:

1. Load the SPD file.
2. Apply the algorithm to the grating positions (and raster points) and produce a new set of scan numbers.
3. The new scan numbers are then written into a new copy of the LSPD file.

This fix to the LSPD product means that the AAL processing stage can then be invoked, allowing the rest of the drift correction algorithm to function properly. Returning to the example, the following plot shows the analysis of LW1 performed after the scan numbering has been correctly written to the LSPD file.

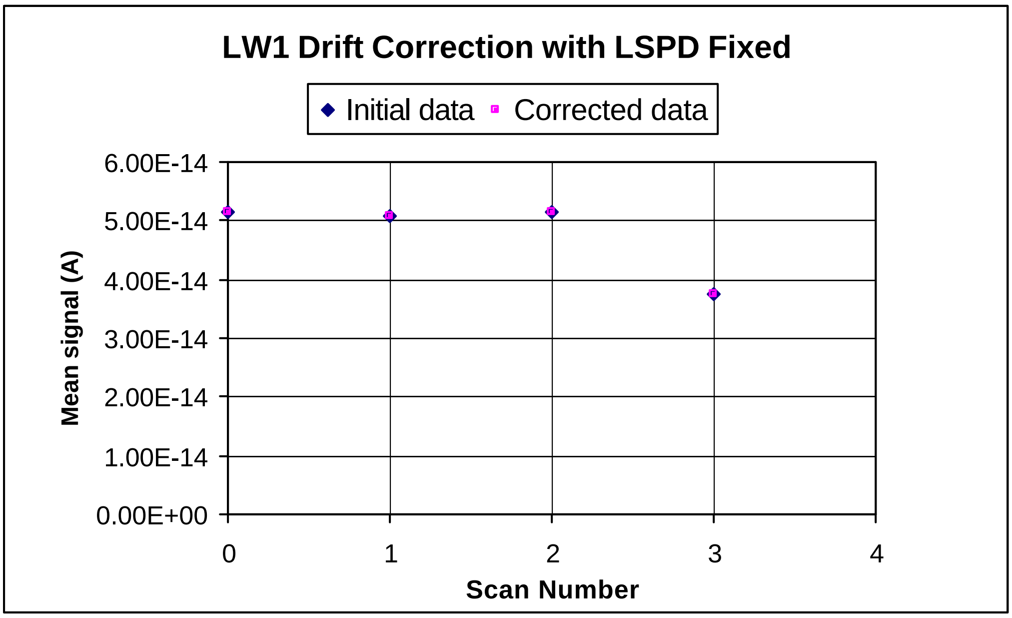

Notice how the initial and corrected values are now very similar, indicating that the drift was indeed minimal on this detector. There are now just 4 data points, since point ‘0’ includes the whole of the first forward scan. Also, the slope fit has correctly ignored the short scan (now numbered ‘3’). The following plot shows what happens to the science data.

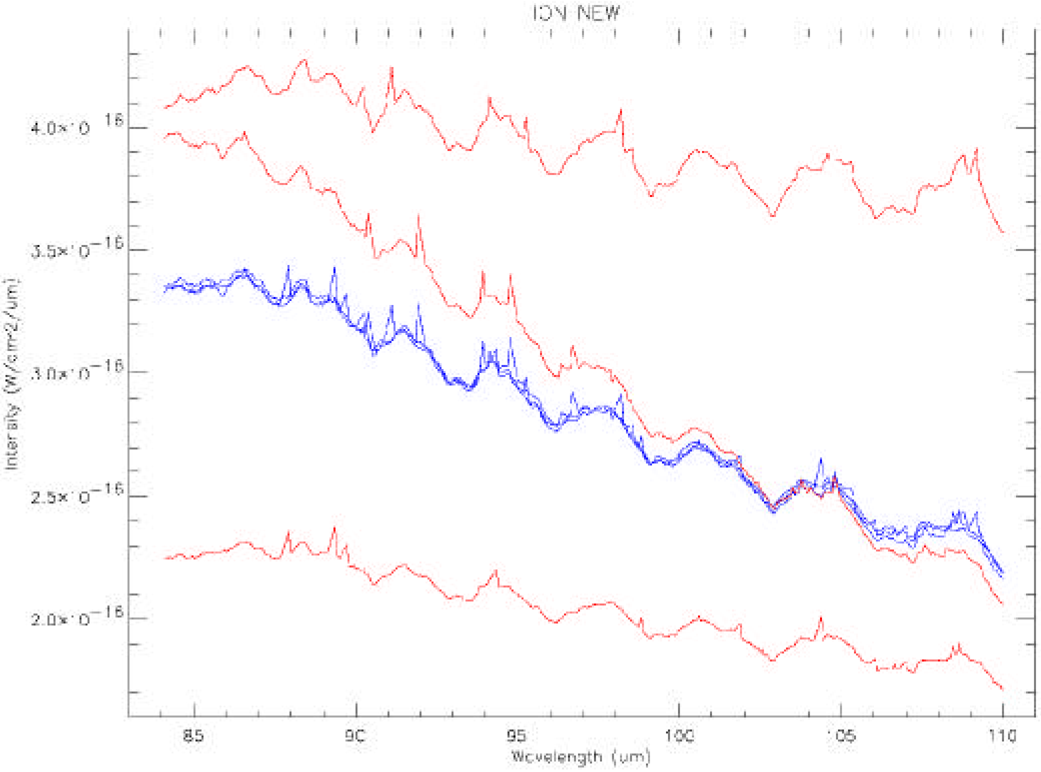

The red lines show the calibrated scans in the original observation, processed with OLP10.1. There appears to be a strong disagreement between scans, in an apparent trend of decreasing detector responsivity. The blue scans show the same data, processed after the scan numbering has been correctly written in the LSPD file. The agreement between scans is significantly improved, and shows the expected result for the drift correction algorithm. The full corrected spectrum is shown below.

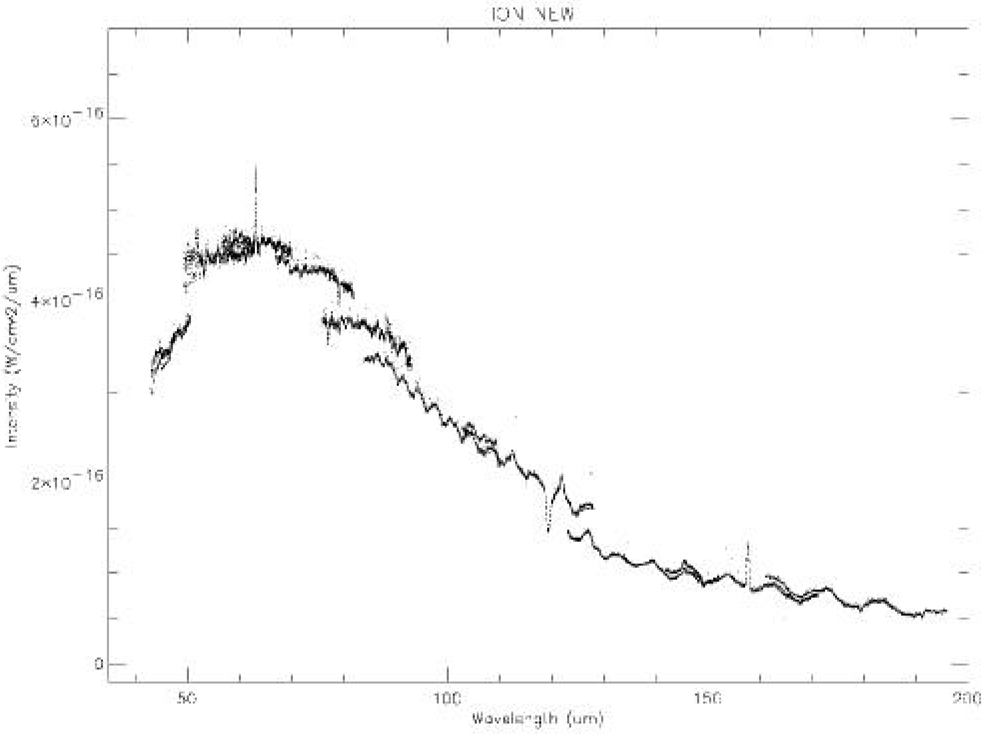

While this observation is not free from other problems (such as strong source effects, transience, and fringing), it is now up to the standard of the rest of the OLP10.1 products.

There are several more algorithms in the AAL, following the drift correction algorithm. They are each applied in turn to the drift-corrected data, but do not depend on the scan numbering in order to function correctly. Therefore, no changes to the AAL were necessary in order to correct the scan-numbering problem. In order to make sure that no other processing elements were affected, several tests were performed in the validation.

Finding and Correcting the Observations Affected

The scan-numbering problem is difficult to find by any direct test. This is because the pipeline result files contain no reference to the intended details of observations, except for wavelength coverage and the active detector used in line AOTs. The intended numbers and directions of scans are not available to the pipeline, so there is no independent means of verifying that the scan numbering is correct. Instead, the best way to look for problems was to examine all the different cases that can occur for LWS01 observations, and check that the algorithm used is capable of giving the correct results. The main criteria that vary are:

1. Number of raster points
2. Number of scans per raster point
3. Number of ramps per grating position
4. Scan direction (normally alternate forward / reverse, but can be forward only)
5. Number of lines scanned (different wavelength ranges)

By manual inspection of observations varying in each of these ways, it was possible to determine that the fixing algorithm was robust under all conditions.

The second stage was then to attempt to apply the algorithm to all LWS01 observations, and observe whether it made any difference to the scan numbers. This was done by comparing the numbering in the existing LSPD files to that generated by the algorithm, and looking for any changes. 277 observations were found to contain 1 or more jumps in scan number, with a total of 310 different raster points affected. A sample of these was inspected by hand to verify the existence of the problem, and all were found to be affected. A sample of products not flagged as affected was also inspected, and all were found to be normal.

Once all of the observations were identified, the full reprocessing of the affected observations was undertaken.

- The OLP10.1 archive versions of the LSPD and LIPD (used in the AAL stage) products were used as input.
- For each observation, the fixing algorithm was applied to the LSPD file, allowing a new set of scan numbers to be derived. This was then written to a new copy of the LSPD file.
- The new version was then used as input to the OLP10.1 auto-analysis pipeline, to produce new LSAN, LSCA and LGIF products. The LIAC products were also produced as usual, but these are the same as those produced before the scan-numbering problem was fixed.

Validation of the Reprocessed Data

The first stage of the validation was to verify that the correction had no effect on other aspects of the calibration procedure. To check this, the LIAC files produced by AAL for the reprocessed products were compared with the original OLP10.1 LIAC files. The dark current and absolute responsivity values for each detector were examined, to ensure that these had not changed. All of the reprocessed LIAC files were checked (with an automated procedure), and none were found to have changed.

The second stage was to check the changes in the drift correction coefficients, and analyse the severity of the effects on each observation. This involved comparing the old and new values of the coefficients, and comparing the magnitude of the drift correction in each case. The drift correction is always largest at the start and end of each calibration group (time periods with different sets of calibration coefficients), as the detector drift is corrected to the responsivity at the centre of the group. Therefore, the values at the edges of each group were used to make the comparison. A ratio was calculated for each detector, comparing the magnitude of the correction with the old (O) and new (N) coefficients. An average was then taken between the values at the start (ST) and end (EN) of the group.

$$Ratio = \left(\left({}^{O_{ST}}\!\big/_{N_{ST}} \right) + \left({}^{N_{EN}}\!\big/_{O_{EN}} \right) \right) \Big/ 2$$

Where e.g. O_{ST} represents the correction factor at the start of the group. The nature of the correction is such that the factors are reversed between the start and end of the group i.e. above 1 at the start, and below 1 at the end or vice versa.

The ratio calculated above could be greater or less than 1, and positive or negative, so a normalisation was then applied to give a positive value greater than 1 for each detector. Averaging over the ratios for all 10 detectors will hence give a value greater than 1, if the group was affected by the scan numbering problem.

The output from this analysis was hence a table with a value for each detector, for each group of each observation. The total number of groups was 379. In LWS01 observations, there can be 1 or more calibration groups per raster point. The following example is an extract from the table. It shows an observation with 1 group per raster, so there is hence only one row per raster.

TDT 31600474

Raster	SW1	SW2	SW3	SW4	SW5	LW1	LW2	LW3	LW4	LW5	Mean
2, 2	1.002	1.018	1.016	1.066	1.023	1.030	1.072	1.044	1.014	1.006	1.0291
1, 3	1.001	1.009	1.010	1.006	1.004	1.016	1.010	1.015	1.000	1.007	1.0078
7, 2	1.003	1.001	1.015	1.003	1.004	1.010	1.014	1.032	1.021	1.024	1.0127

The table was then further analysed, to give the average effect on each detector per raster point. This makes sense, since science data is averaged over entire raster points, so the effects in 1 group compared to another will be smoothed over in the analysis. A second table was hence produced, listing the fractional difference in drift correction for each detector in each raster point; there are 310 raster points in total. This gives an indication of the error in the original product compared to the reprocessed product, in terms of the maximum fractional error in the flux in each detector. It also provides an easy lookup for finding out how badly any part of any dataset was affected.

The third stage of validation was to determine the extent of the science impact of the problem. The values calculated in the previous stage were averaged over each detector, to find the observations most severely affected in average terms. This allowed the raster points to be put into rank order, from the most to the least severely affected. The observation used as the example earlier, TDT 84900612, is the second most severely affected observation overall, with an average flux error of 38.4% at the start and end of the observation.

To evaluate what different average errors values indicate, a sample of datasets was inspected across the whole range. In general, observations with average errors higher than around 8% show serious effects on the science quality. This can include changes in continuum slope, inter-detector calibration and the noise in the spectrum. The following plot shows a comparison of the LW1 detector from TDT 84900612, before and after correction.

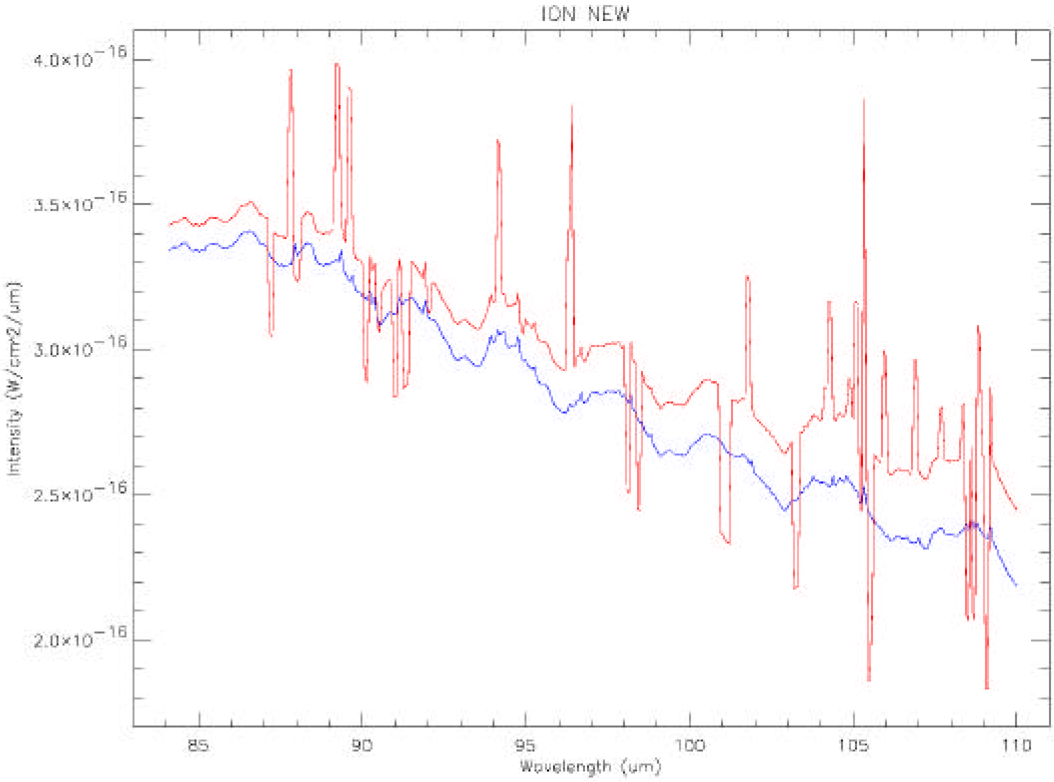

The red line shows the averaged spectrum before correction, and the blue line is the reprocessed data after correction. It is clear that all three effects are present in this detector, with an apparent change in absolute calibration and continuum slope. The noise is not real noise, but an artificial noise caused by glitches. A glitch in one of the scans will mean that the average of the two remaining scans is changed significantly, as the scans themselves are not well matched. When the drift correction functions properly, this type of effect is not visible. In general, this effect is more prominent in observations with low numbers of scans, and for sources with a high signal-to-noise ratio in the underlying continuum.

After examining the full range of average errors, the following scheme was determined, in order to rate the severity of effects in different raster points.

Quality Flag	Extent of Effects	Description
1	Trivial	No overall science impact is likely, although minor effects may be present in some detectors
2	Minor	Overall science may change slightly; moderate effects possible in some detectors.
3	Moderate	Overall science probably affected; serious effects possible in some detectors.
4	Serious	Significant change in overall science likely; severe effects possible in some detectors.

The full table with these flags assigned to the raster points has been made available to the ISO data centre. These should allow users to determine whether OLP10.1 data that they have already reduced is still valid, or if the reprocessed products should be used instead. In general, it is likely that users of datasets rated ‘3’ or ‘4’ will have to use the reprocessed products. Approximately 110 raster points fall into one of these two categories. Datasets rated ‘1’ or ‘2’ will probably not be too bad overall, but users should note the warning that some detectors could be worse than others. For example, in the table shown on page 8, raster (2, 2) has an average error of 2.9%, which puts it in category 2. However, detectors SW4 and LW2 have errors of 6.6% and 7.2% respectively, so are more likely to show some science impact.

As individual detectors can show significant variation from the general behaviour in an observation, even datasets rated ‘1’ or ‘2’ should not be used for data reduction. It is therefore proposed that new users be provided with only the reprocessed products, and hence the original products in the archive should be replaced.

Conclusions

A scan-numbering problem has been identified in 277 LWS01 mode observations, out of a total of 1791 fully processed datasets. A correction method was devised and implemented to fix the problem at SPD level. This was tested on all products to verify that it correctly handles all possible different LWS01 mode settings. The correction method does not affect the science in the LSPD products, and LIPD products are unchanged. All observations affected have been corrected at SPD level.

AAR products were reproduced from the 277 sets of corrected SPD products, using the OLP10.1 auto-analysis pipeline. This resulted in changes to all LSAN, LSCA and LGIF products from these observations. LIAC products were also reproduced, but were confirmed to be unaffected by the correction. As the scientific algorithms used did not require changes, the scientific quality of the new products is that documented in the OLP10.1 validation report.

A sample of the reprocessed products was checked by hand, and all were found to have been correctly produced. A systematic validation was then carried-out on all the reprocessed products, to examine the differences between these and the original OLP10.1 versions. This allowed the severity of the effects on the science to be rated, giving an indication of the extent to which the original products fell below OLP10.1 scientific standards.