The UVES Data Reduction Pipeline

P. BALLESTER[1], A. MODIGLIANI[1], O. BOITQUIN[2], S. CRISTIANI[1], R. HANUSCHIK[1], A. KAUFER[1], S. WOLF[3]

[1]European Southern Observatory; [2]Royal Observatory of Belgium, Brussels; [3]Science Computing GmbH, Munich

1. Introduction

The Ultraviolet-Visible Echelle Spectrograph (UVES) was the first instrument installed at the second VLT Unit Telescope KUEYEN at the end of 1999. UVES is a high-resolution cross-dispersed echelle spectrograph covering the wavelength domain from the ultraviolet at 300 nm to the near-infrared at 1100 nm with two optical arms, which can be used in dichroic mode for parallel observations (Dekker et al., 2000). Its large-size echelle gratings, each 84 cm long, make it possible to achieve a spectral resolution of up to 80,000 in the blue arm and 115,000 in the red. It is equipped with 3 CCD detectors each of size 4K × 2K pixels, one EEV CCD in the blue arm, and an EEV/MIT CCD mosaic in the red arm. Commissioning observations have shown the high intrinsic efficiency of this instrument, particularly in the ultraviolet domain (13% at 360 nm). More than 200 hours of scientific observations from UVES commissioning and science verification time have been released for public use. The instrument started regular operation in visitor and service mode on April1, 2000, and is used for about 75% of the telescope time (D'Odorico, 2000).

The need to provide quick-look reduction of the data at the telescope, to monitor the instrument performance and to prepare the data distributed to service mode observers are the major drivers for VLT instrument pipelines. In the case of UVES, with the 500 spectral formats supported by the instrument and the expected 2.5 Gbytes of daily data, special emphasis was put into designing methods for automated spectral format calibration supported by physical models. These methods were prepared in 1998 and the pipeline integration started in 1999. After the initial phases of integration in Europe and commissioning of the instrument, the pipeline was first installed at the telescope during the second commissioning of the instrument in December 1999. It was then used for technical verification, and during the Science Verification that took place in February 2000. It was upgraded in March 2000 with a robust version, improved optimal extraction, support for image slicer mode, an improved background estimation procedure and a complete update of the calibration database. The pipeline was then adapted to meet the needs of the Data Flow Operations (DFO) group in Garching, which delivers reduced data to service mode observers. A further upgrade with increased checks of quality control parameters was done in August 2000. Extended-source extraction, as well as the extension to support the FLAMES fibre port of UVES are under development.

2. Pipeline Operations

The UVES pipeline is tuned primarily to support operations in Paranal and Garching. The main difference between Paranal and Garching use of the pipeline lies mostly in the operational objective of the reduction. In Paranal it is used to have a quick look at the on-going observations and monitor the instrument health. DFO operations aim at providing the best-possible reduced products and at controlling their quality. As a consequence, the update rate of the calibration solutions differs between the two sites: daily in Garching, monthly at Paranal.

The Paranal pipeline is the one available to visiting observers at the telescope. It supports all standard settings of the instrument (11 central wavelengths and 4 detector modes) and makes use of pre-determined calibration solutions established for a reference configuration. The accuracy of this calibration is therefore limited by possible instrumental drifts for example due to temperature variations or earthquakes. The data-reduction software is robust to shifts of up to about ten pixels, and the spectral format is regularly compared at the telescope with reference calibration solutions. If a deviation is noticed that goes beyond the limit of robustness of the pipeline, it is necessary either to realign the instrument to the reference configuration or to update the database. Non-standard settings are also supported at the telescope. These solutions are however not stored in the calibration database, but in a temporary area and are removed at the end of an observing run.

The Garching pipeline is the one used by the Quality Control Scientists to prepare the reduced data distributed to service mode observers. The best possible accuracy is achieved drawing from the pool of available calibration data – either produced during the observing run, or as part of regular calibration programmes – the ones judged most suitable by the DFO Quality Control Scientist. The Garching pipeline applies the same pre-defined reduction strategy to all science data, independent of the scientific purpose of the observations. The quality of the calibration solutions is also measured and stored for trend analysis.

3. Reduction Procedures

To work its way through the many configurations, the pipeline makes use not only of the FITS header information, but also of the specific knowledge of the optical design of the instrument by means of a physical model. The model is involved at many stages of the calibration to provide reference values and initial solutions for the current configuration making the data reduction completely automatic. In addition to the calibration solutions, the pipeline delivers quality control information to help the user in assessing the proper execution of the data reduction process.

The UVES pipeline is able to provide good-quality science-data reduction for low to medium signal-to-noise data under the assumption that the object is a point-source, centred on the slit. The standard spectra extraction in the course of the reduction at Paranal or in Garching is made in optimal extraction mode which provides a maximal signal-to-noise ratio per bin, automatic sky background correction and cosmic ray rejection. Different reduction strategies can be implemented by using directly the data reduction package, which provides additional extraction methods and data reduction options.

3.1 Predictive format calibration

The geometric calibration is a complete definition of the spectral format including the order position and the

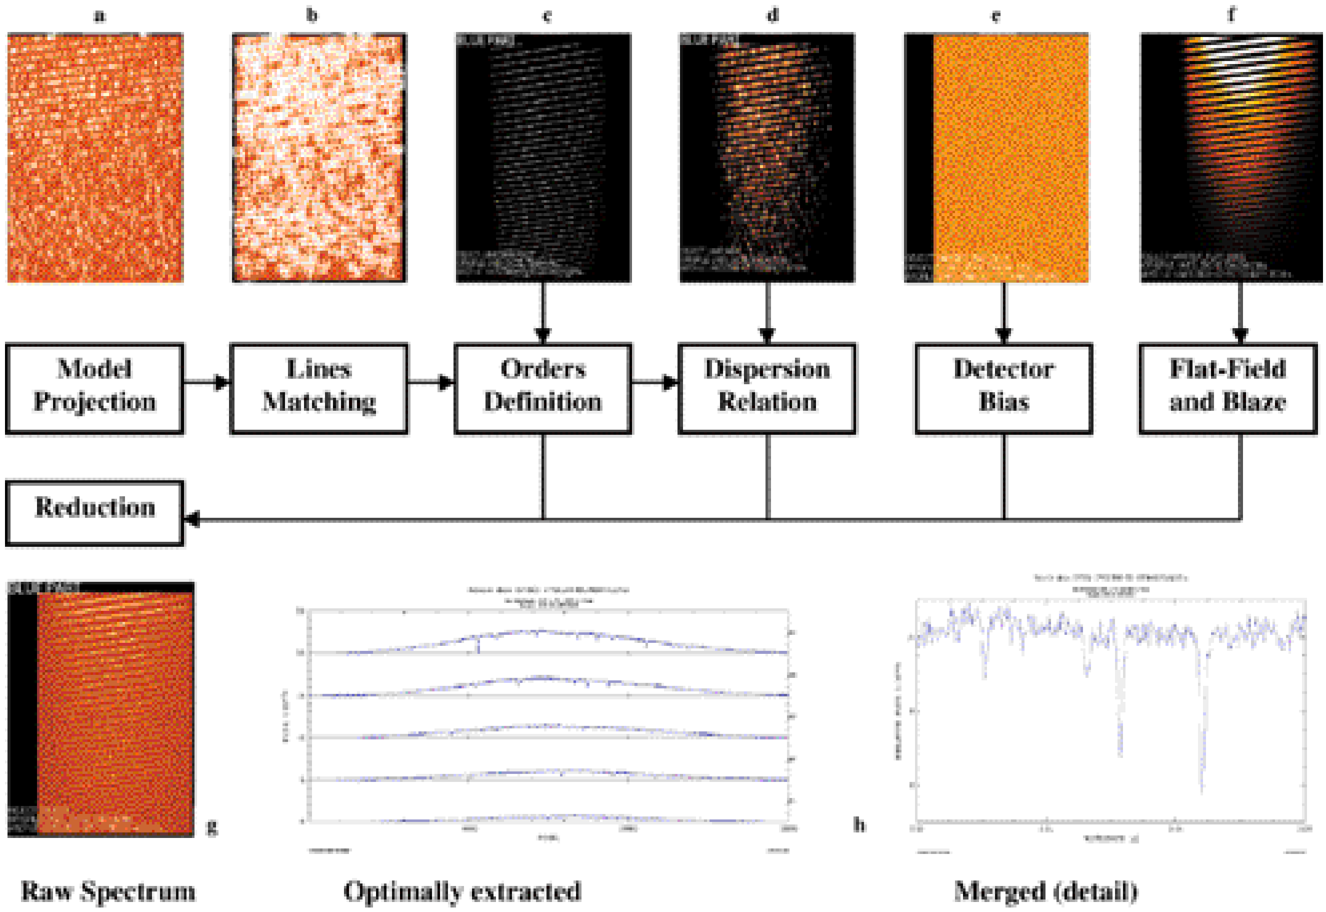

Figure 1: The general reduction scheme of the UVES pipeline. Model predicted positions of wavelength calibration lines are projected onto a format-check frame (a). The predicted positions are adjusted to the observed positions and an initial dispersion relation solution is produced (b). The order position is automatically defined on a narrow flat-field exposure using the physical model and a Hough Transform procedure (c). The initial dispersion relation is refined on a Th-Ar frame in order to fully take into account the slit curvature (d). Bias exposures are averaged to produce a master bias frame (e). Flat-field exposures are processed to produce a master flat frame and to determine the blaze function (f). The science frames are automatically reduced using the previous calibration solutions (g). After optimal extraction and order merging, a one-dimensional spectrum is produced (h).

wavelength associated to each detector pixel. This step was traditionally carried out via visual identification of a few lines and for this reason many new methods had to be developed. The precision with which the geometric calibration is performed determines the accuracy of all successive steps, in particular the optimal extraction and the photometric calibration. Initial methods for an automated detection of the order position and the wavelength calibration were developed using the Hough Transform method (Ballester, 1994). In the case of UVES, the high non-linearity of the dispersion relation required the development of a physical model of dispersion in order to predict the dispersion relation and to efficiently calibrate the several hundreds of optical configurations (125 central orders in 4 detector modes) offered by the instrument (Ballester & Rosa, 1997).

A template is dedicated to acquire a "format-check" frame (Fig. 1a, b), namely a ThAr calibration taken with a short slit height. The positions of a few hundred well-separated ThAr lines contained in a reference table are predicted by the physical model and their central positions are projected onto the format-check frame. The lines are found in the narrow Th-Ar frame by a two-dimensional centring procedure. The initial dispersion relation, usually based on about a hundred initial detections is refined with successive iterations on the spectrum until most lines are found.

The format-check data-reduction step generates control plots which provide a quality check of the model and of the stability of the instrument configuration. The difference between the predicted and measured position of the reference calibration lamp lines are plotted as a function of the position on the detector or the wavelength. Typically, a thin, well-clustered distribution of the displayed points, with mean value near zero, is a direct indication of good model prediction (Fig. 2a). On the contrary, a randomly scattered distribution indicates a probable instrument configuration shift (Fig. 2b). In this case one could guess adjusted model parameters. Once a concentrated distribution is found again, usually by changing the offset along the cross-order direction, the plot will have a mean value close to the actual shift of the instrument configuration in the corresponding direction (Fig 2c). After the May 12, 2000 earthquake event, the instrument control parameters have been readjusted to a reference configuration.

A preliminary step toward the reduction of echelle spectra is the definition of the order position (Fig. 1c). For the sake of accuracy and stability of the slit tracing, this operation is usually performed on contrasted, continuous spectra such as flat-fields or bright star exposures. A specific narrow flat-field template was defined for this purpose. It takes an exposure of the flat-field lamp with the shortest slit equal to 0.5 arcsec. This kind of exposure provides an accurate position of the spectrum along the cross-dispersion direction. In this step the physical model uses the information on the instrument configuration provided in the FITS header to estimate the number of orders present in the image. Hough Transform detection is then applied to find the central position of each order and an estimate of their slope at the centre. Finally the cross-order profile is centred along the order and a polynomial fit is performed.

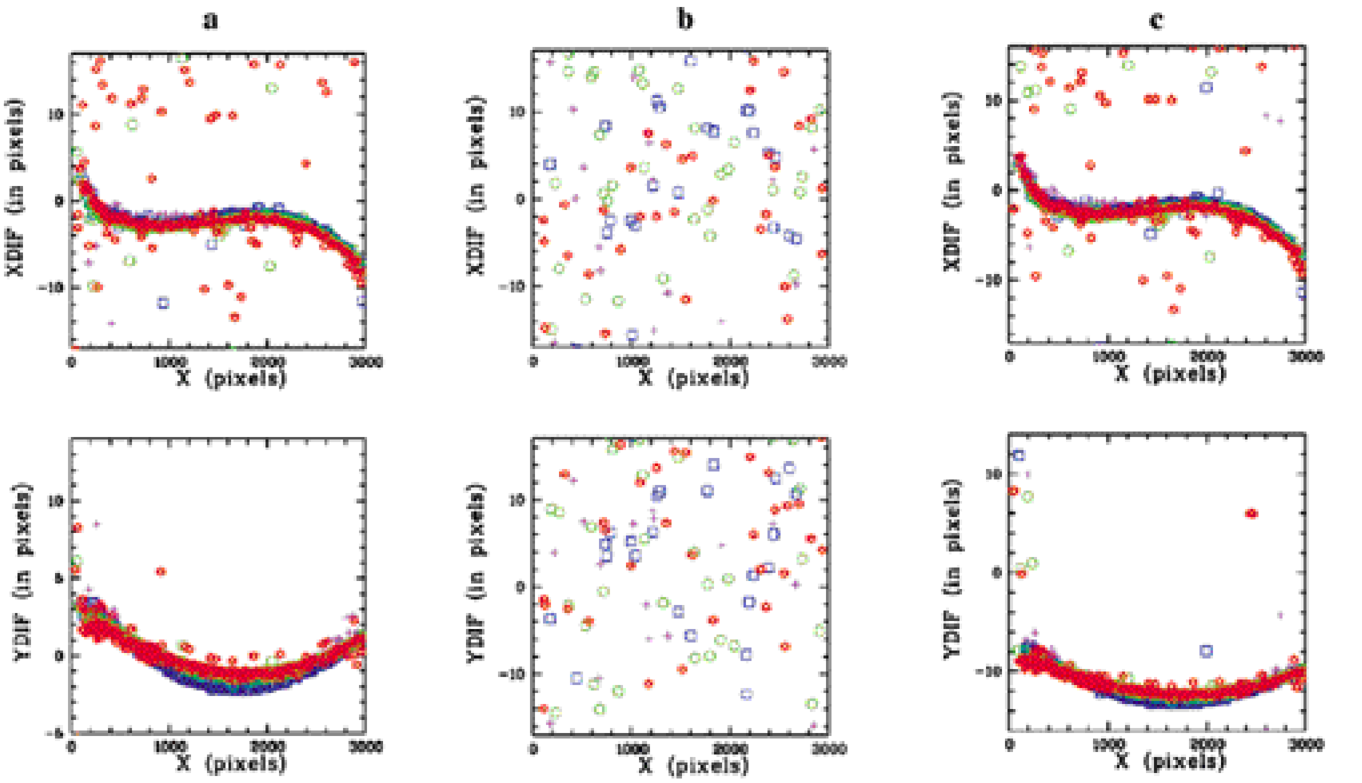

Figure 2: The May 12, 2000 earthquake event as detected from the physical model control plots. The normal result obtained after success - ful line matching (a) produces a well concentrated distribution with mean ordinate zero. The earthquake event causes the lines matching step to fail (b). Adjusting the model by –10 pixels (along the cross-order direction) again matches the instrument configuration (c).

Three independent wavelength calibrations (Fig. 1d) are performed along the slit to take into account the potential effects of slit curvature. Three sub-windows are identified along the slit, one for the object and two for the sky, all equal in size. The initial dispersion relation produced with the format-check frame is refined on each sub-window.

From a set of bias exposures it is possible to produce a master bias image through a stacked mean (Fig. 1e). Similarly from a set of flat images it is possible to create a master flat image with master bias subtraction and background subtraction as well as to derive the blaze efficiency function for each order (Fig. 1f). Science raw data (Fig. 1g) can be processed using optimal or average extraction.

3.2 Optimal and average extraction

Optimal extraction is a minimal variance estimator of the flux at each wavelength. The algorithm introduced by Horne (1986) for long-slit data assumes that the illumination fractions of the spatial profile vary smoothly with the wavelength and can be adjusted by low-order polynomials. This assumption does not generally hold in echelle spectroscopy due to the inclination of the orders. Resampling the data along the spatial profiles would introduce periodic profile distortions and noise. Different methods have been developed for cross-dispersed spectroscopy (Marsh, 1989; Mukai, 1990) involving data resampling for the sole purpose of estimating the weights.

The optimal extraction algorithm of the UVES pipeline is based on profile fitting by chi-square minimisation. It estimates the object signal and sky background and performs cosmic ray rejection under the assumption of a Gaussian profile of the light distribution along the slit. The extraction is performed independently on each order. First the full-width at half-maximum and the position of the object is estimated for different subsections along the order. Next through a chi-square minimisation, the amplitude and the sky background are estimated. The Gaussian shape is a reasonable assumption for low to medium signal-to-noise data. Sky emission lines are removed during the extraction: they fill the whole slit and are accounted for by the constant term of the spatial profile. Tests have shown that the sky fine subtraction achieves an accuracy better than 5%. Strong sky emission lines may remain visible in the extracted spectrum as residual light and of course by the larger variance of the extracted spectrum. Cosmic rays are rejected by comparing the raw cross-order profile with the result of the fit. One of the common difficulties in optimal extraction schemes is to prevent the rejection of valid data samples by the cosmic ray rejection methods in particular on high signal-to-noise data. In the UVES pipeline, the rejection threshold is therefore adjusted to the signal-to-noise so that the cosmic-ray rejection is relaxed for increasing signal-to-noise data, converging to an average extraction scheme for high signal-to-noise data. Presently, this method appears to be appropriate for data with a signal-to-noise per bin up to fifty. For higher signal-to-noise, an average extraction is recommended which is also provided in the data reduction software.

In the average extraction, the source is centred on the slit and sky windows are determined. The object contribution is summed in the object sub-window, and the sky contribution is averaged and subtracted from the object. Data acquired in image slicer mode are processed with average extraction and without sky subtraction.

Both extraction methods, average and optimal, support the propagation of variance. The initial variance is estimated on the raw images by a noise model including read-out and photon noise. The variance images are then transformed together with the data frames on each processing step. Finally this information is used to optimally merge the orders and to deliver error bar estimates on the result spectrum.

3.3 Standard pipeline processing and general reduction

The standard pipeline processing of science data starts with data preparation: the raw input frames are rotated

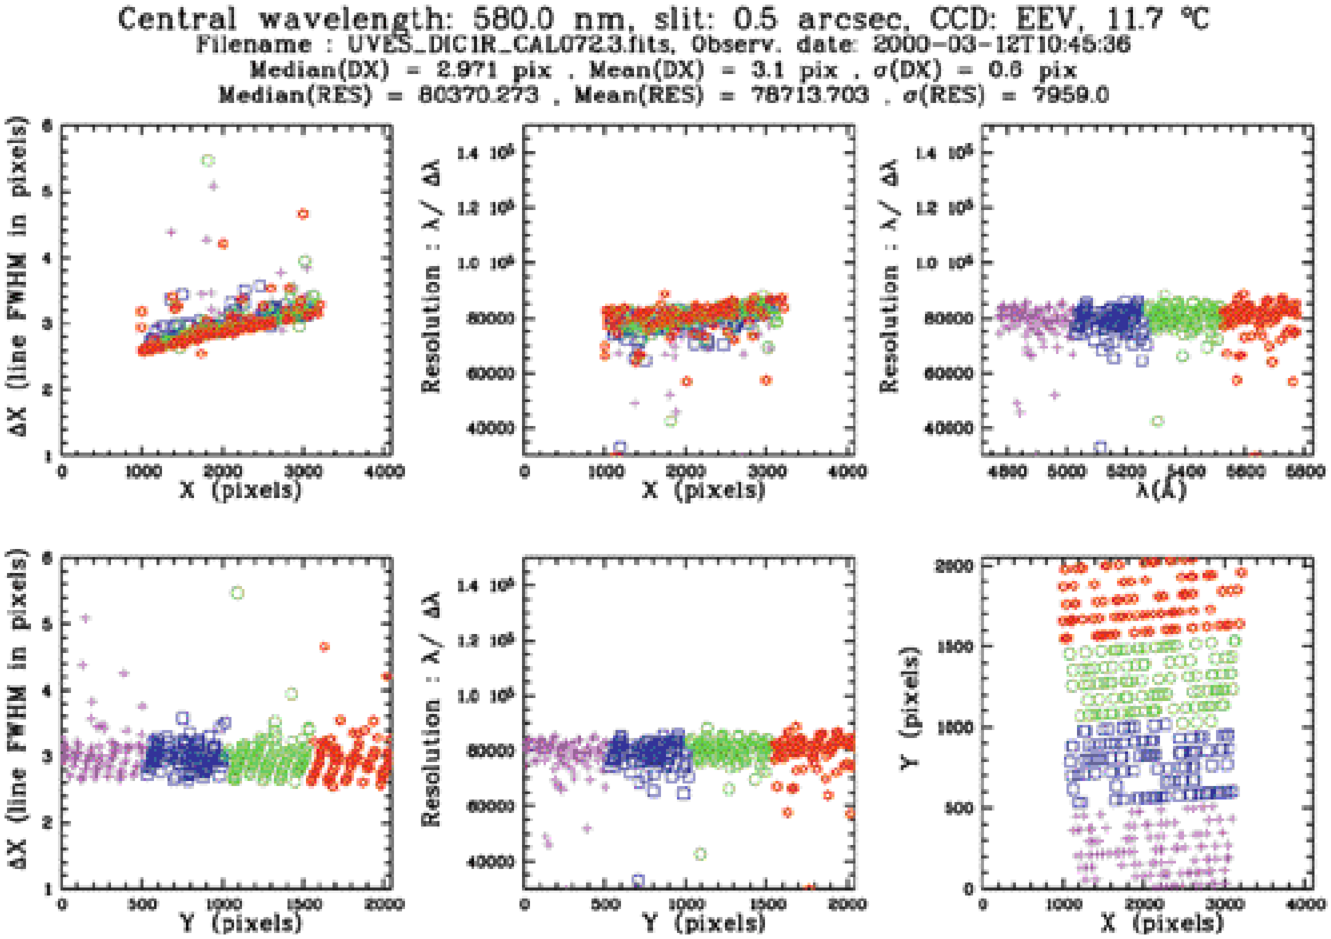

Figure 3: These plots show the FWHM of ThAr lines (left) and the instrument resolution (centre) as a function of the position on the detec - tor along the X and Y directions, and finally the resolution versus wavelength (upper right) and the distribution of measured lines over the detector (bottom right).

and in the case of the red mosaic data, split for each detector. The instrument's background light is corrected in two steps: a master bias frame is subtracted and the inter-order background light is estimated by a minimum filter. The optimal extraction provides the object signal and its variance. The flat-field correction is performed in the standard reduction scheme in pixel-order space on the extracted data. The signal is then resampled to constant wavelength bins and the orders are merged into a single spectrum. The most time-consuming steps in the science reduction process are the inter-order background subtraction and the optimal extraction of the spectra. On Paranal, a 4K × 2K science exposure produced by one red detector is fully reduced in about six to twelve minutes depending on the machine.

In echelle spectroscopy, for a proper merging of the adjacent orders, a prerequisite is the correction of the blaze function.The UVES pipeline makes possible to estimate the blaze function from a spectrophotometry standard or from a flat-field exposure. The analysis of standard stars makes use of the target coordinates to look up in the database the appropriate flux reference table, reduces the star spectrum and determines the conversion factor for each order. In standard pipeline reduction, the blaze function is estimated from a flat-field exposure, which after extraction, wavelength calibration and smoothing enables merging of adjacent orders with an accuracy of a few per cent.

The data reduction software and pipeline procedures are distributed as a MIDAS context on the 99NOV CD-ROM. It includes all the pipeline procedures for data reduction and on-line quality control, documentation, step-by-step examples and a tutorial demonstrating the complete calibration and reduction process for blue arm data. With this package, any user can interactively reduce the data, changing the options or the data selection with respect to the observatory pipeline. The user can interactively produce for all UVES configurations the calibration solutions necessary to properly reduce the raw science data.

4. Pipeline Quality Control Procedures

Almost as important as data reduction is quality control. A number of procedures have been written to verify the instrument performance in terms of stability, efficiency and resolution. With format-check exposures a stability check can be performed. Using the physical model, the line position on the actual format-check frame and on a reference format-check frame are measured and then the mean and standard deviation of the relative difference are determined. This quality control procedure has been particularly important to track the stability of the spectral formats since the beginning of operations (Figs. 2, 4).

During wavelength calibration, the mean, median, and standard deviation of line FWHM, and of the spectral resolving power are monitored as a function of the position on the detector. The resolution plot (Fig. 3) shows information on the instrument spectral resolution derived during the process of wavelength calibration.

On Paranal, the UVES pipeline records the quality control parameters in log files. Spectral resolution, instrument stability and efficiency parameters are tracked and stored together with environmental parameters measured by UVES sensors. These log files are archived for long-term trend analysis, and parsed for the generation of graphs accessible within the observatory to the operations teams. Figure 4 shows the evolution of the frame shift along the dispersion as a function of the observing date. This variation is correlated with the evolution of the temperature in

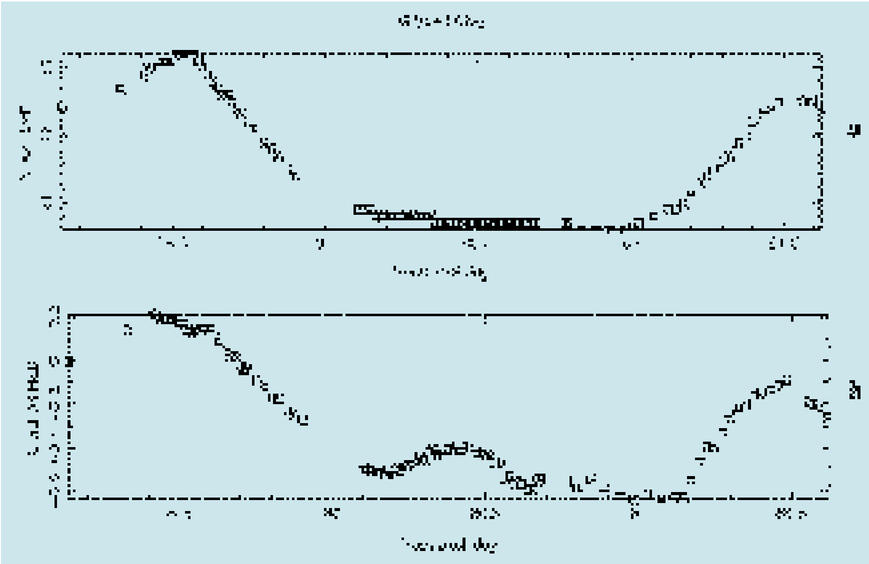

Figure 4: Evolution of the temperature in the instrument enclosure (a) and average shift in the dispersion direction (b) as a function of time over a period of 2.5 days. The two graphs show a strong correlation.

the instrument enclosure. The QC logs provide highly valuable information to the observatory to monitor the performance of the instrument and, if necessary, to recover or improve its quality.

The efficiency for each order at the blaze wavelength is monitored. The overall efficiency of the observation system, consisting of the telescope, instrument and detector, is determined as a function of wavelength by reducing spectrophotometry standards for which a flux table is available. It delivers an efficiency curve which makes it possible to directly track the reflectivity of optics coatings in the different wavelength domains and to verify the instrument performance predicted by the Exposure Time Calculators (http://www.eso.org/observing/etc).

The procedures for CCD characterisation are based on the MIDAS CCD package. Bias and flat-field exposures with different illumination levels are analysed to verify the CCD sensitivity and linearity maps. Usually a set of twenty-four exposures, six biases and nine pairs of screen flats exposures, each pair taken with the same exposure time, are used to determine several CCD parameters. For the 1×1 binned flat-field exposures, the intensities are measured in nine reference sub-windows, uniformly distributed on the detector to detect possible edge contamination. One field is taken as reference. The relative intensities with respect to this reference field are logged for quality control.

5. Quality Checks by QC Garching

Although planned as an automatic data reduction 'machine', the UVES pipeline is also used by the Quality Control (QC) Scientist in Garching in a semi-interactive way to derive best-possible calibration solutions and for reducing science data obtained in Service Mode. Procedures checking on the quality of pipeline results are essential here. They have been developed as post-pipeline procedures by the UVES QC Scientist to evaluate many of the pipeline-generated results.

As an example, Figure 5 shows the evolution of the spectral format in cross-dispersion direction (Y) for the blue (top) and red (bottom) gratings. These shifts are determined from flat-field exposures against a reference flat frame. They clearly correlate with UVES temperature. In addition, the May 12, 2000 earthquake has shifted the detector of the blue arm by more than 10 pixels.

Figure 6 shows a typical example of how the quality of a reduced (blue) spectrum is assessed.

1. The central row and column of the raw exposure are plotted. Here one can monitor any anomaly of the detector signal, e.g. unusual bias level.
2. A close-up of a central column (actually of 20 averaged columns) shows the (instrumental plus object) profile in cross dispersion direction. Here the exposure level, the centring and the sky background are controlled. These parameters are essential for assessing the quality of the optimum extraction.
3. Lower-left diagram: the full spectrum (S), its variance (N) and the signal-to-noise ratio (S/N) are plotted. The extreme left ends of S and N are shown at two different scales to fit in the graphic window. These

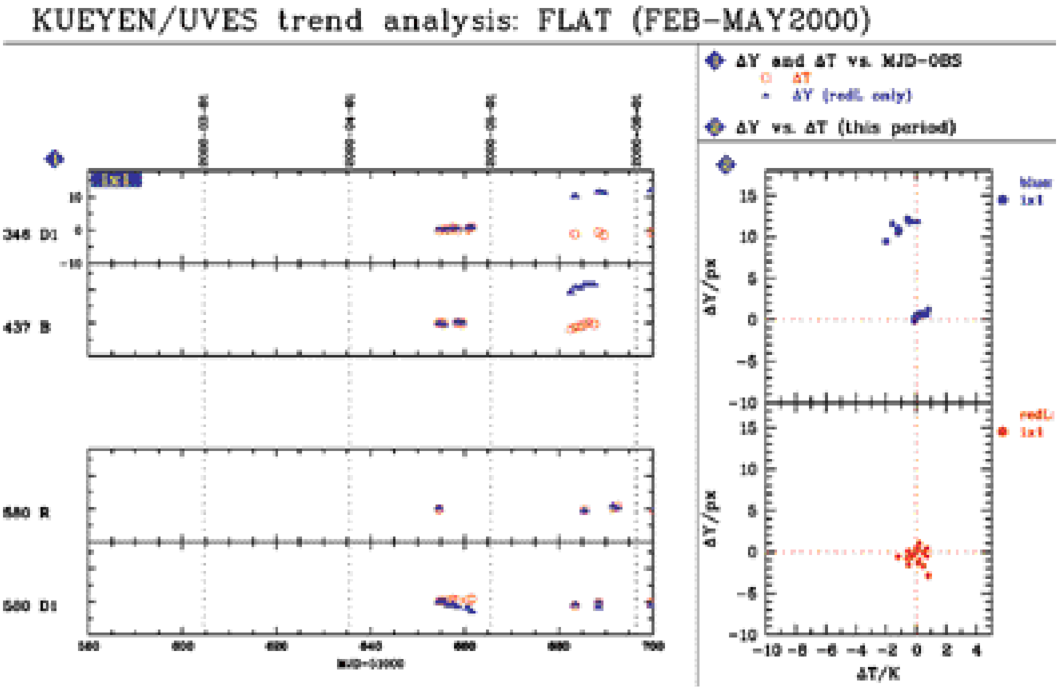

Figure 5: Left-hand-side panels: Vertical (cross-dispersion) shifts ΔY against temperature ΔT for the blue (top) and red (bottom) gratings. For four of the 2 × 11 standard settings, ΔY (blue triangles) and ΔT (red circles) are plotted over time to monitor their trending. Right-hand-side panels: A composite diagram includes the two boxes on the right to visualise the thermal drifts of all gratings. The sudden non-thermal shift by more than 10 pixels on the detector of the blue arm was caused by the May 12, 2000 earthquake.

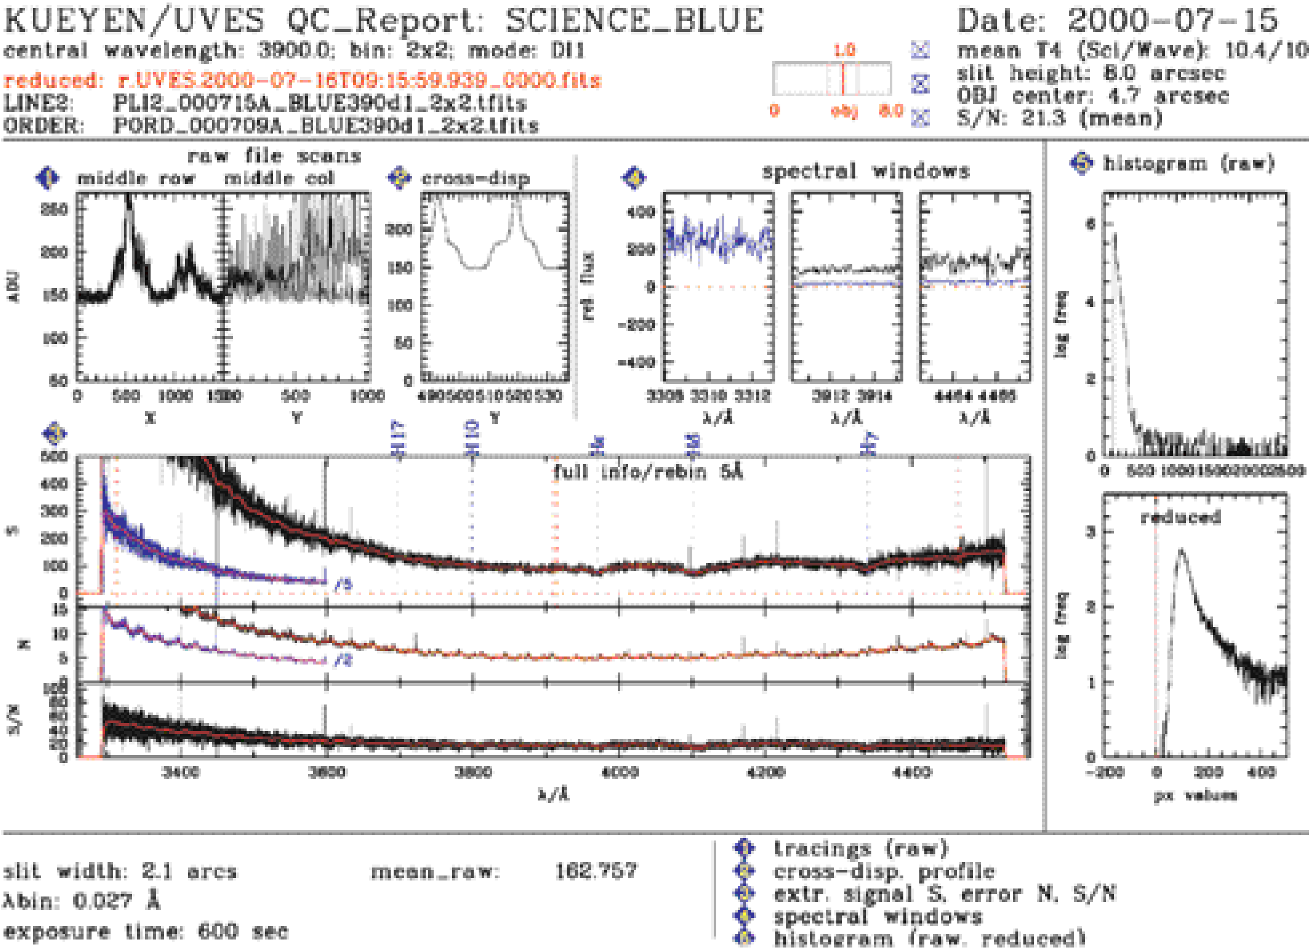

Figure 6: QC plot for a blue reduced science frame. See text in Section 5 for details of the plot.

tracings permit an immediate check on the quality of the extraction.

4. Three spectral windows show enlarged portions of the full spectrum to check single bins.
5. Histograms of the raw and the reduced data provide further checks on anomalies.

Finally, three check-boxes (Fig. 6, upper right) raise alarm flags when

- the temperature difference between science and ThAr calibration frame is larger than a threshold value,
- the object is closer than 20% to the ends of the slit (check on centring). To the left of the middle check-box is a sketch of the slit showing the object position and the seeing FWHM.
- the signal-to-noise ratio is higher than a threshold value (then optimum extraction becomes unsafe).

Currently these plots are used to build up a trending database and to gain expertise for the quality assessment. It is planned to deliver these plots to the Service Mode observers. The UVES results from Quality Control Garching will be soon made public on the web (http://www.eso.org/observing/dfo/quality/).

During the first four months of operations, from April 1st to July 31st, 2000, more than 14,000 UVES raw files were produced representing about 153 Gbytes of data and including more than 4000 science frames. About half of these frames have been taken in service mode, resulting in 2640 reduced science spectra and 3852 calibration files produced with the pipeline.

Conclusion

The UVES pipeline illustrates the increased efficiency and stability to be gained by the systematic use of instrument models and variance propagation in the reduction process. The cost of establishing new calibration solutions has been dramatically reduced, so that quick-look data reduction at the telescope is supported for most of the several hundreds instrument configurations. The instrument performance is monitored by comparing the calibration results with respect to the model and by measuring the effect of environmental parameters (temperature, earthquakes). All standard configurations of the instrument are supported in service mode and data packages are regularly shipped. A quality control database is being established for trend analysis.

Acknowledgements

Many thanks to the UVES Principal Investigator, S. D'Odorico, the project responsible, H. Dekker, and the UVES optical designer, B. Delabre, who provided many essential contributions to the design of this package since the beginning of the project in 1995.

References

Ballester, P., 1994, Hough Transform for robust regression and automated detection, *Astron. Astrophys.*, **286**, 1011–1018.

Ballester, P., Rosa, M.R., 1997, Modeling echelle spectrographs, *Astron. Astrophys. Suppl.*, **126**, 563–571.

Dekker, H., D'Odorico, S., Kaufer, A., Delabre, B., and Kotzlowski, H., 2000, Design, construction and operation at the telescope of UVES, the Echelle spectrograph for the UT2 Kueyen Telescope at the ESO Paranal Observatory, Proc. Conf. SPIE 4008-61.

D'Odorico, S., 2000, UVES at Kueyen: Bright prospects for high-resolution Spectroscopy at the VLT, *The Messenger* **99**, March 2000, 2.

Horne, K., 1986, An optimal extraction for CCD spectroscopy, *PASP*, **98**, 609–617.

Marsh, T.R., 1989, The extraction of highly distorted spectra, *PASP*, **101**, 1032–1037.

Mukai, K., 1990, Optimal extraction of cross-dispersed spectra, *PASP*, **102**, 183–189.