

Figure 7: The upper panel shows a medium resolution ($R \sim 3000$) spectral image around $1.06 \mu\text{m}$ of two galaxies at $z \sim 0.62$ selected from the CFRS catalogue. Total exposure time was 1 hour. The lower panel shows the extracted spectrum of the brighter galaxy which, in addition to the continuum, shows bright redshifted $H\alpha$ and fainter [NII] (to the right) emission lines. As can be seen in the image, the $H\alpha$ emission appears to be tilted relative to the continuum due to the rotation of the galaxy which can thus be measured to provide an estimate of its mass.

7. Acknowledgements

We wish to particularly thank the ISAAC Science Team R. Chini, G. Miley, E. Oliva and J.-L. Puget for their suggestions and encouragement throughout the design and development phase; R. Gilmozzi, P. Gray, M. Petr, G. Rahmer and P. Sans-

gasset for their support during installation and commissioning on Paranal and J. Alves, P. Amico and M. Kissler-Patig in Garching. Collective thanks are also due to many other people, too numerous to mention individually and within all Divisions of ESO, who have contributed to this project. We are also grateful to L.




Figure 8: A spectral image of SN 1987A in the HeI ($1.083 \mu\text{m}$) line obtained in a 20-min exposure using the medium-resolution grating and a $2''$ wide slit. The ring is formed of gas blown off the progenitor star during its evolution and which was first photoionised by UV light from the SN and more recently shock ionised by ejecta which, travelling at $30,000 \text{ km/s}$, has only just reached it after a travel time of nearly 12 years. The jet-like feature extending to the left is actually in the spectral dispersion direction and shows the presence of a blue-shifted, high-velocity component that is believed to arise in the ionised ejecta.

Tresse and O. Le Fèvre for detailed information on the CFRS galaxy sample.

References

- Hammer, F., Crampton, D., Le Fèvre, O., Lilly, S.J. 1995, *ApJ*, **455**,88.
 Lizon, J.-L. 1996, *The Messenger*, **86**, 11.
 Meyer, M., Finger, G., Mehrgan, H., Stegmeier, J., Moorwood, A.F.M.: 1996, *The Messenger*, **86**, 14.
 Moorwood, A.F.M. 1996, In *Optical Telescopes of Today and Tomorrow* (ed. A. Ardeberg), *SPIE* Vol. 2871, 1146.

amoor@eso.org

ISAAC Pipeline Data Reduction

N. DEVILLARD, Y. JUNG, J.-G. CUBY, ESO

The ISAAC pipeline, just as the pipeline for any other VLT instrument, is built to provide reduced data on site in order to assess their quality as soon as possible. Another version of the same software will also be running in Garching for quality control, trend analysis, and service-observing data-reduction purposes. All instrument pipelines share the same generic structure, taking care of data handling, i.e. transfers and reduction recipe triggering. The recipes themselves are specific to each instrument, they are plug-ins hooked into this generic data-handling structure, and they can be described as specialised tasks attached to an instrument-acquisition template. For more details about pipelines and Data Flow, the reader is referred to the ESO Web pages

and corresponding reference documents.

The following instrumental modes are supported in pipeline mode for the first ISAAC period: imaging jitter, twilight flat, darks and bad pixel map creation, illumination frames and zero-point computations. In spectroscopy: NodOnSlit, flat-field, spectroscopic response function, star trace calibration, and slit position. We give here an overview of the algorithms used to reduce the data without user intervention for all these modes, so that observers can get a rough idea of what is done in pipeline mode, what the working domain for the recipes is, what can be obtained in the control-room at runtime, and how to reproduce these reduction processes at their home institute if they wish to.

1. Imaging Mode

Imaging recipes (procedures) are based on the ESO **eclipse** library, mainly to speed up the number-crunching process (see article in *The Messenger* No. 87, p. 19–20). The eclipse library is exportable, compiles on virtually any (recent) Unix machine that has minimum POSIX compatibility, and is available on the Web at the following address: <http://www.eso.org/eclipse/>

Calibration recipes

Dark frames are acquired through a dedicated template, which stores several dark frames for each possible DIT (Detector Integration Time), and usually

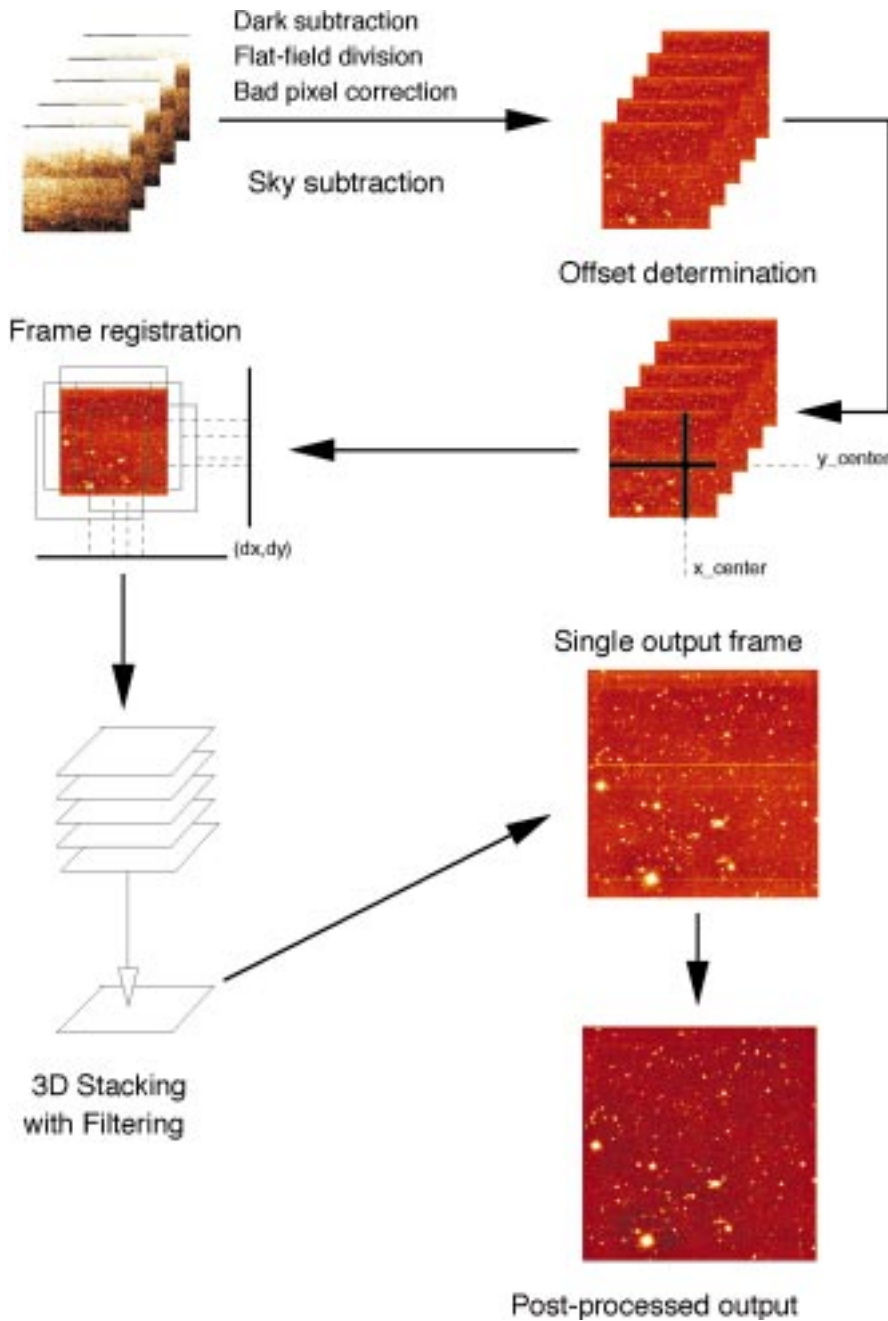


Figure 1.

covers several DIT values. The output of this template is a list of frames containing dark exposures, which need to be sorted out by DIT group and then averaged per group with rejection criteria to avoid cosmoics. The output of the **is_darks** procedure is a list of frames, one for each given DIT. Read-out noise is also measured from these frames and the information is written to a readable ASCII file. The associated eclipse routine is **is_darks**. It is recommended to take 5 darks or more per DIT you used during the observation night, so that the computed master dark frame is reasonably reliable.

Twilight flatfields are acquired through a template that observes the sky at dawn or dusk with a constant integration time. Produced frames have an increasing or decreasing signal value that is used to measure detector linearity. For each pixel on the detector, a curve is plotted of the

median plane value (x axis) against the individual pixel value in this plane (y axis). This curve shows the pixel's response, from which a robust linear regression provides a pixel gain. The image showing all pixel gains (i.e. the flatfield) is normalised to have an average value of 1. Some criteria can be added later to check that the pixel behaviour is indeed linear, and build a bad pixel map for pixels that have excessively high or low gains, or do not behave linearly at all. The same procedure could be used with any kind of increasing/decreasing signal, which allows to sample the pixels' response at various points. The associated eclipse routine is called **flat**.

Flat-fields are usually corrected for low-frequency variations which are not taken into account by the twilight flat procedure. To compute these variations, a dedicated calibration template is offered. The procedure images a standard star on a grid

of 4×4 or 5×5 positions, with the convention that the first acquired frame shall contain the standard at the centre of the image, and subsequent positions be indicated through keywords in the FITS header. Photometry is computed for each star position and a 2d second-degree polynomial is fitted through the obtained surface. An image is generated from this polynomial with an average value of 1. The associated eclipse routine is **is_illum**.

Zero-point acquisitions are done by observing a standard star with wide offsets on the detector. The operator at the instrument is in charge of locating the standard star through finding charts and positioning it at the centre of the array for the very first frame. Subsequent frame offsets are indicated in the FITS header. The recipe identifies the exact position of the standard star up to ± 50 pixels from the centre, computes the number of counts, and relates that measurement to a standard star database. The possible telescope pointing errors are not relevant to this case, because the star has been identified and positioned by human means. This is necessary to ensure that the star is actually in all frames and this is part of the acquisition process for this template. As a consequence, the algorithm does not need automatic astrometric calibration but a simple local peak-finder method. The ideal would be to provide star localisation through un-supervised astrometry; this may be implemented later. The identified infrared star database at this moment contains about 800 star positions with magnitudes in bands J, H, K, Ks, L, M. Some stars have missing information, it is recommended to test out beforehand that the standard star you will be using can actually be found and used from this database. The star search only looks for the closest star in position from RA and DEC information in the header, and computes zero-point information from it. In case of star mismatch (e.g. several standards are in the field), it should be easy to check out that the star used corresponds to the acquired one because all data relative to the used star are given along with the measured values. The standard star database is currently accessible through an eclipse command called **std**. Later on, access should be given through the Web to this list and users should be able to provide their own standards. The associated eclipse routine for zero-point computation is called **is_zerop**. Notice that this method will not apply to standard stars in crowded fields, since it is relying on the assumption that the standard star to look for is the brightest object in the centre of the first frame, a condition that is not fulfilled in that case.

Observation recipes

The main instrumental mode supported by the pipeline is the jittering mode (also known as microscanning). In this mode, no direct sky observation is performed, which maximises observation efficiency. However, a number of as-

sumptions must be fulfilled in order to use this template. The observation procedure is to first centre the first frame on the target, then acquire images with small telescope offsets between frames. The offsets requested to the telescope are stored along in the FITS headers. Of course, the usable part of the data seen through the detector is the intersection of all frames so the offsets should not be too large. On the other hand, the very aim of these offsets is to show what the sky background is to the detector, so offsets should be larger than the largest object on the frame. This means that the method can only be used with relatively small objects, which is typically the case for deep fields. For extended objects or crowded fields, the option is to apply this jittering to two positions: one on the target and another one on an empty field anywhere nearby, the immediate result being a loss of observing efficiency without which no sky subtraction is possible. In this case, we are talking about a jitter+offset mode. The idea in that case is to use the jittered pattern on the (empty) sky frames to filter out all possible remaining object signals and produce a clear sky that is subtracted from object frames. For the moment, the method running in the pipeline is taking a median average of all sky frames and subtracting it from the objects, but that would only work for small numbers of frames. Later implementations will offer various sky filtering algorithms depending on the actual acquisition duration and filter used.

In pure jitter mode (all frames acquired on the object), the instrument has a built-in offset generation method that scatters points optimally with a Poisson distribution law. The generated offsets have the property to be optimally distributed over the possible range of offsets, and thus maximise the chances of being able to filter out the sky background (see the `poisson` command in eclipse). It is recommended to use the `AutoJitter` template unless you have specific needs for the offset values and provide your own to the `GenericOffset` template. Please note that this generic template cannot yet be supported in pipeline mode, because it leaves too many choices to the user to be able to support all possibilities. When we get experience about what is usually requested, some more reduction modes will be implemented for the most usual cases.

The reduction procedure for the jitter algorithm is shown in Figure 1. Pre-processing is done by applying dark subtraction, flatfield division, and bad-pixel correction if these calibration frames are available. The first algorithm step is then to filter out the sky background from raw images. This is done with a 3d-filter with pixel rejection, actually similar to the `imcombine` routine in IRAF. Sky statistics are accumulated at that point and written to an ASCII result file; they should be appended to the output file FITS header in later versions. The second step is to detect the offsets between frames. We




Figure 2.

need some useful signal in the frames to do that, and the best configuration is to be able to find the same bright object in all consecutive frames. If this is not possible, the algorithm would look for all bright objects it can find and try to locate them image after image, with sub-pixel precision. This is done through cross-correlation methods and is usually quite robust and insensitive to noise. The third step is to register all images to the first image taken as a reference. This is done through re-sampling methods in image space, and produces much less aliasing than Fourier space-based methods, linear interpolation, or just integer pixel moves. The only visible consequence is that noise (high frequencies) is smoothed out by the re-sampling process, but photometry is conserved. Once this is done, the last step is to stack all frames to a single one, using again a 3d-filter to remove spurious values. Post-processing is applied to this image to remove detector row saturation effects for example, then the result may be sent to an image displayer.

Figure 2 shows an image produced in November 1998 with ISAAC, the Galaxy cluster CL2244-02 with gravitational arcs (see press release dated 26 November 1998 on the ESO web site, and the December 1998 *Messenger* issue). This image is the jittered combination of 60 input frames in Ks band; it was obtained with default settings for all parameters. This is a typical output from the ISAAC pipeline in the control-room, for that kind of data.

This algorithm proves to be very reli-

able and usually complete for typical fields acquired so far with SOFI or ISAAC. There are still some incorrectly processed cases though, mostly for images where all bright objects are located on the edges of the first frame (the reference), but that should be solved with e.g. point-pattern matching algorithms in later releases of the pipeline. For more detailed information about the algorithm and what can be fine-tuned, readers are referred to the following document: "jitter imaging data reduction: algorithms and implementation".

The URL is <http://www.eso.org/projects/dfs/papers/jitter98/>

The software running automatically on both SOFI and ISAAC pipelines takes about 25 minutes (on a fast workstation with large amounts of memory) to reduce a batch of 60 frames, i.e. roughly 250 megs of data in input, and can run on virtually any POSIX-compatible machine. Supported OS's so far are Solaris, HP-UX, Linux, IRIX, AIX, DEC OSF/1, and any station that offers POSIX compatibility and an ANSI C compiler.

2. Spectroscopy Mode

The spectroscopy pipeline is made of MIDAS procedures making use of the LONG and ECHELLE packages, with some calls to eclipse routines to speed up pure image processing routines, and to perform standard star database search. It is likely that we will redevelop in C identified parts of the algorithm that are not likely to change and need CPU speed.

Calibrations

Various calibration procedures are performed for trend analysis purposes; they are mainly used by the instrument scientist. They include a control of the slit position with respect to the detector, calibration of the spectroscopic dispersion relation with calibration lamps, analysis of known spectral lines to correct for spectrum tilts or curvature in the spatial direction, and spectroscopic flat-fields.

Observations

Spectroscopic observations performed with the *NodOnSlit* template produce spectra on two different positions A and B obtained by a shift of the telescope along the slit. Each set of images is averaged to a single image, and both sets are subtracted from each other. Available outputs are images of (A-B), (B-A), and a combined image of those two. The com-

bined image is actually the most important one for the observer, since extraction will take place on this 2d-spectrum image. The pipeline goes on from this point, extracting the brightest spectrum from this image, performs flux and wavelength calibrations when adequate calibration data are available, and outputs the results to a FITS table.

The brightest spectrum is to be seen as a checkpoint: since it is usually a reference object, it might be useful to check out that the acquisition actually performed as expected at least on this object. This spectrum is available to the user in the control room, together with the combined 2d image so that spectrum extraction can proceed from that point (partly) by hand.

Notice that spectroscopy recipes should be distributed as a part of MIDAS in a subset dedicated to pipeline routines. Some are supporting an interactive mode that actually allows giving more indica-

tions about how the procedure runs and lead them to proceed directly to the solution for the spectra of interest.

Acknowledgements

All ISAAC pipeline data reduction algorithms are based on the ISAAC calibration plan (ESO document VLT-PLA-ESO-14100-1384) by Pierre-Alain Duc and Jean-Gabriel Cuby. All procedures have been improved tremendously following the advice from Chris Lidman. Numerous discussions with Alan Moorwood, Almudena Prieto, Paola Amico, Pascal Ballester have helped a lot to clarify the requirements and improve the spectroscopy parts. The jitter process has received constructive and helpful comments from too many people to quote them all here; we thank them for their contributions.

ndevil@eso.org

Commissioning FEROS, the New High-resolution Spectrograph at La Silla

A. KAUFER, O. STAHL, S. TUBBESING, Landessternwarte Heidelberg-Königstuhl
P. NØRREGAARD, Astronomical Observatory Copenhagen
G. AVILA, P. FRANCOIS, L. PASQUINI, A. PIZZELLA, ESO

Introduction

On November 30, 1998, the second commissioning phase of the new Fiberfed Extended Range Optical Spectrograph (FEROS) was completed at the ESO 1.52-m telescope at La Silla.

The instrument had been installed by a small team from the Heidelberg, Copenhagen and La Silla observatories starting in mid-September – just two years after the contract signature between the FEROS consortium and ESO. FEROS saw its first stellar light at the end of the installation phase on October 6th.

An overview of the design and of the expected capabilities of the FEROS instrument has been presented in a previous paper [Kaufer et al., 1997, *The Messenger* 89, 1]; a more detailed description of the opto-mechanical design is found in [Kaufer & Pasquini, 1998, Proc. SPIE Vol. 3355, p. 844]. In the present article, the major technical results from the two commissioning phases which followed the first light event are reported.

The Spectrograph Environment

Figure 1 shows the FEROS spectrograph inside its temperature and humidity controlled room, which was erected by the La Silla infrastructure team at the place of the former Echelec spectrograph inside the coudé room of the ESO 1.52-m telescope. The feet of the optical

bench are sitting on a platform coupled to the telescope pier but decoupled from the room's floor. The LN₂ supply tank is located outside of the spectrograph room in the neighbouring FEROS maintenance room and can be replaced every two weeks without entering the spectrograph room. Adding to these facts the very stable temperature which is maintained inside the room ($\Delta T < 0.2$ K over 1 night), FEROS is located in a very stable environment – a crucial prerequisite for a long-term stability of the instrument as required for high-precision spectroscopic measurements.

From the spectrograph, the two fibers together with their metallic protection cable are led through the telescope's polar and declination axes, further inside the telescope's double tube, and eventually leave the telescope and enter the Boller & Chivens (B&C) Cassegrain spectro-

graph to end in the telescope's focal plane, which now hosts the new FEROS fiber head and the previous B&C long slit.

Instrument Performance

Table 1 summarises some of the main performances of FEROS as determined during the commissioning.

Efficiency

The detection efficiency of the complete system ESO 1.52-m telescope + fiber link + FEROS + Detector has been measured several times under photometric conditions using bright and faint CTIO standard stars [Hamuy et al., 1994, *PASP*, 106, 566]. Correction for the atmospheric extinction has been made according to the standard La Silla extinction tables. The typically achieved detection

Table 1: Measured performance of FEROS.

Wavelength Range in one exposure (object+sky)	360 – 920 nm (39 orders, 2 fibers)
Resolving Power	$\lambda/\Delta\lambda = 48,000$ (2.2 pixels)
Entrance Aperture	2.7 arcsec (circular)
Detection Efficiency (with telescope and detector)	1% at 360 nm, 16% at 440 nm, 17% at 550 nm 16% at 640 nm, 11% at 790 nm
Limiting Magnitudes at the ESO 1.52	16 mag in V, S/N = 9, $t_{\text{exp}} = 2$ h 12 mag in V, S/N = 105, $t_{\text{exp}} = 2$ h
Radial-Velocity Accuracy	21 m/s rms (2-month time base)
Inter-order straylight	< 3%