Phoenix: automatic science processing of ESO-VLT data

Reinhard Hanuschik*[1]
*European Southern Observatory, Karl-Schwarzschild-Str. 2, D-85748 Garching, Germany

ABSTRACT

ESO has implemented a process to automatically create science-grade data products and offer them to the scientific community, ready for scientific analysis. This process, called *'phoenix'*, is built on two main concepts: 1. a certification procedure for pipelines which includes a code review and, if necessary, upgrade; and 2. a certification procedure for calibrations which are processed into master calibrations, scored and trended. These master calibrations contain all information about the intrinsic instrumental variations and instabilities inevitable for ground-based telescopes. The *phoenix* process then automatically processes all science data using the certified pipeline and the certified master calibrations. *Phoenix* currently focuses on spectroscopic data. The first *phoenix* project has been the processing of all science data from UVES, ESO's high-resolution Echelle spectrograph at the VLT. More than 100,000 Echelle spectra of point sources, from begin of operations (March 2000) until now, have been reduced and are available to the public from the ESO archive, `http://archive.eso.org/cms/eso-data/eso-data-products.html`. The *phoenix* process will also feed future UVES data into the archive. The second project has been X-SHOOTER slit spectroscopy which currently has more than 30,000 Echelle spectra from the UV to the infrared (up to 2.5µm). The *phoenix* process will be extended to other, mostly spectroscopic, instruments with certified pipelines, like FLAMES. Also, all future VLT instruments will be supported by *phoenix*.

Keywords: Science-grade data products, pipelines, data standards, data processing, quality control

1. INTRODUCTION

In ground-based astronomy, a night observation of a scientific target contains information from three different sources: the scientific target (more precisely: any extra-terrestrial source along the selected line of sight), the terrestrial atmosphere, and the instrument itself (telescope mirrors, optical elements inside the instrument, detectors). While generally these contributors cannot be disentangled on the measurement ("raw data") itself, it is true that at least the disturbing components can be measured separately. The signature of the atmosphere can be recorded with standard stars, and the signature of the instrument can be registered with standard light sources and detector exposures. These calibration data have to be taken in due time and in appropriate frequency. Their acquisition pattern forms the *calibration plan*.

At the Very Large Telescope (VLT) of the European Southern Observatory (ESO) on Paranal/Chile, each of the currently 15 instruments has its own calibration plan. The science data taken during the night trigger the appropriate set of calibration data taken during the day after, whenever possible. For some calibrations, precious night time has to be used, namely for standard stars for photometry, telluric absorption etc.

If done properly, these calibrations register all instrumental and atmospheric properties correctly, completely, and appropriately. This third important condition means that the calibration data need to be applicable to the science data, and they are applicable only if the conditions they should measure (e.g. encoder positions, temperatures, efficiencies, and many more) are the same during daytime, when calibrations are taken, and night time. This is one of the most important requirements on the acceptance of science data taken at the VLT: they must be calibratable, otherwise they need to be repeated.

With the calibrations done properly, the science data can be reduced, meaning the instrumental and atmospheric signature can be removed from the science data to a known level.

[1] rhanusch@eso.org; phone 49 89 3200 6742; www.eso.org

At ESO, the Quality Control Group is processing all VLT calibration data. The products get their main properties measured (with "QC1 parameters") which are trended (put into the context of other similar QC1 parameters) and scored (assessed in terms of compliance with pre-defined thresholds) [1]. With this information evaluated, and with expert knowledge available if something unusual occurs, the calibration products (the master calibrations) are finally certified (accepted as complying with the specifications), or rejected. A certified master calibration can be seen as a valid measurement of the prevailing instrumental (or atmospheric) conditions that can be applied for the reduction of science data.

The certified master calibrations are stored in the ESO archive.

2. PIPELINES

The VLT is organized as "science data machinery", with any data acquisition standardized. Data are acquired in observing blocks and templates, are recognized by the Data Flow System (DFS) and processed by the pipelines. These pipelines are usually delivered by the consortia building the instruments and then integrated at ESO into the DFS.

The pipelines have their focus on the removal of instrument signature. They are used both at Paranal, for quick quality control and quick-look data reduction during the night, and at ESO Headquarters in Garching bei München for final quality control.

Experience shows that pipelines are not always complete, stable and accurate. Over the years they can be improved via bug reports and fixing mechanisms. Some pipelines (e.g. UVES and X-SHOOTER) have been in operations for a couple of years and then saw a major review including a review of the algorithms. Such pipelines are now well understood and stable and can be called 'certified'.

3. SCIENCE-GRADE DATA PRODUCTS

For a certified pipeline, the adopted reduction strategy, as well as the error budget, are well known and understood. With such a pipeline it is possible to reduce data towards the *science-grade level*, defined as a science product with all relevant instrumental and atmospheric effects removed to a well-understood level. Also, for science-grade data products a reduction strategy has been chosen that is independent of the science case. This independence is naturally enforced by the fact that the science case cannot be automatically derived from the headers, and also from the strategic goal that at that processing level, a data product can be used for the analysis of a science case that is different from the original one. Such kind of re-use could be called 'archive science'.

The science-grade, or science-ready, data products are the ultimate goal of a pipeline-based data reduction, in contrast to the publication-ready, or 'advanced', data products that have some knowledge applied about the physical nature of their targets or their original scientific goal. For instance, knowing the exact position of your source and its physical properties, you might be able to extract a faint signal on top of a complex background, something that requires a customized (meaning, most likely, interactive) strategy which goes way beyond a standard, source-independent strategy, being automatic and 'simple'.

For at least two kinds of research, the science-grade data products offer big advantages: for the above-mentioned archive science, and for the scientists browsing the archive for statistical projects, or for already existing observations of targets of their own field of interest. In all those cases a science-ready data product is likely to be preferable over a set of raw files requiring the complete processing chain including knowledge of many different data types and processing parameters: *per aspera ad astra*. Even for the PI, the automatically reduced science-grade data products may be attractive since they can serve as reference or even as acceptable output from where the scientific analysis can start.

With the combination of certified pipelines and certified master calibrations, the reduction of ESO data into science-grade data products has become possible. Since all instrumental effects are recorded properly in the master calibrations, the science data products, even at a massive scale, can be created automatically and reliably. This is a non-trivial achievement for ground-based spectrographs, due to all the possible instrument interventions, issues and fixes which make the data quality measurement and control quite complex.

In ESO terminology we call these data products the Internal Data Products (IDPs) since they are processed in-house. We also offer External Data Products (EDPs) provided by the PIs of the Public Surveys (at the VST and VISTA survey telescopes) [2], [3]. The EDPs are typically *advanced data products* being processed into catalogues and mosaics. Being based on surveys, they represent a very different kind of science, as compared to the IDPs which reflect the integral over the spectroscopic interests of the entire ESO community over the last 14 years (see section 5 for some numbers).

4. PHOENIX: SCIENCE-GRADE DATA PRODUCTS FOR SPECTROSCOPY

In the following we will focus on the IDPs. We have designed a workflow tool for the process for creating them, called '*phoenix'*. It has been applied to the data delivered by the two main spectroscopic workhorse instruments on Paranal, UVES and X-SHOOTER.

With all certifications been done, no decisions need to be taken for the processing of science data. The strategy (translated into pipeline processing parameters) is standard and documented. The association (the selection of the proper master calibrations to go with the science data) is simple in most cases (closest-in-time), and controlled in the exceptional cases by breakpoints (introduced by a sudden change of instrumental conditions, e.g. a detector change, or by an earthquake).

The *phoenix* process is executed by a shell script tool with the same name that controls the whole workflow. Here is a sketch of the *phoenix* process for a typical batch of the UVES data:

1. The tool is called for a given full night.
2. The associations (list of master calibrations required for a given science dataset) are downloaded from a central repository.
3. The raw science data are downloaded from the archive.
4. The master calibrations are downloaded from the archive.
5. The processing jobs use the science recipe of the pipeline and are launched in parallel (for efficiency) on a multi-core machine, using a scheduling engine.
6. Quality-control jobs are scheduled and executed in a similar manner. They extract quality parameters and prepare preview plots.
7. The data products are ingested into the archive.

For the production of science-ready data products at ESO, we have decided to focus on spectroscopic instruments which combine a potentially high interest for the community and a substantial gain for the user since spectroscopic data reduction is generally more complex than the reduction of imaging data.

5. THE UVES ECHELLE DATA PRODUCTS

UVES is the optical high-resolution Echelle spectrograph at the VLT, offering a spectral resolution of up to 110,000 between 320 and 1000 nm. It started operations in March 2000 and is still being operated. Its main mode is single-source slit spectroscopy, with optional additional optical elements like image slicers and an absorption cell. For all data being acquired in this mode, the *phoenix* tool has created an IDP. Data for extended objects are not covered by this project, and the multi-object FLAMES/UVES spectra are not included either.

In total, we have processed about 100,000 IDPs for the period between March 2000 and April 2014, with new data becoming available every month. They are available for selection and download from the ESO archive interface for spectroscopic data products[2].

The products come in the 1D spectroscopic data product standard[3], which defines a binary table format with columns for the wavelength, flux, error etc. Find a detailed description of the data products and the data reduction in the release description[4]. There is also a documentation page about the 1D format[5].

---

[2] http://archive.eso.org/wdb/wdb/adp/phase3_spectral/form . Select 'any' for `Programme`, and 'UVES_ECHELLE' for `collection`.

[3] Defined in http://www.eso.org/sci/observing/phase3/p3sdpstd.pdf , search for the parts about spectroscopic data.

As of April 2014, the UVES IDPs come from 1109 different run IDs, based on 909 different science cases (corresponding to 909 proposals[6]). Since UVES can observe in a blue and a red arm simultaneously, typically (but not always) one blue and one red spectrum belong together, they are stored separately. Many observations consist of two shots in separate wavelength setups, in order to cover the entire spectral range offered by UVES. There is always one raw spectrum processed into one product spectrum.

Technically, the processing of one Echelle spectrum takes typically one minute or less, due to the mature and efficient pipeline. The processing was done on a 32-core Dell-M820, with 30 cores configured for parallel execution by *HTCondor*[7]. This setup is efficient enough to effectively process the entire UVES Echelle archive in about one week, with roughly half of the time spent on downloading the raw science and master calibration data into the system, and uploading the IDPs to the archive.

The pipeline recipe consists mainly of the steps:

- the bias level and structure is removed;

- the science data are divided by the flat-field;

- they are extracted order by order using the order table ('pixel' method) and optimal extraction (sky background and cosmics are removed);

- the extracted orders are wavelength-calibrated (using the line tables), rebinned, and merged;

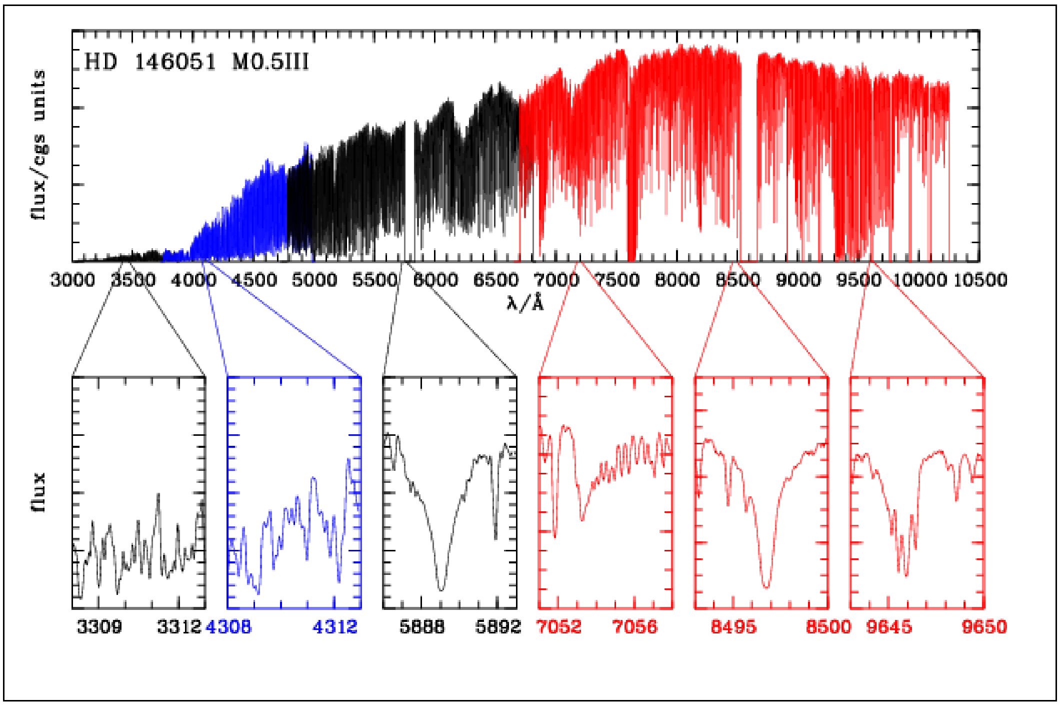

Figure 1. Example of UVES IDPs produced by *phoenix*. This is the fluxed optical spectrum of an M0.5III giant, composed of 4 individual UVES spectra ranging from 300nm to 1μ. The spectral components are coded black-blue-black-red from left (in different grey tones in the printed version). Some close-ups of spectral features are displayed at full resolution. The SNR is about 100-200, the "noise" in the upper panel is actually spectral structure. These spectra have been reduced in an entirely automated way.

---

4 Release description: http://www.eso.org/observing/dfo/quality/PHOENIX/UVES_ECH/processing.html

5 http://archive.eso.org/cms/eso-data/help/1dspectra.html

6 In the QC database http://archive.eso.org/qc1/qc1_cgi?action=qc1_browse_table&table=uves_science_public, you can select run IDs and browse for the abstracts (if already public) of the proposals associated to the UVES IDPs, in the column 'abstract'.

7 *HTCondor* is open-source scheduling software developed at the University of Wisconsin-Michigan.

- if taken in one of the standard setups (a standardized set of central wavelengths), the data are finally flux-calibrated using a master response curve and come in physical units; if not, the final spectrum is not flux-calibrated.

The exact knowledge about the spectral format (i.e. the positions of the Echelle orders, the spectral dispersion within the orders) which is required at high accuracy to preserve the full information of the high-resolution spectra is completely contained in the master calibrations. In the case of UVES these are taken every day (the most critical ones: arclamp calibrations, spectral format calibrations), or every third day (flat fields). Since these calibrations have all been reviewed and certified in the previous day-to-day QC process, the processing of the science data can build on those and be executed in a completely automatic way. See Figure 1 for an example spectrum that consists of 4 individual spectra (2 pairs of spectra from the blue and red arm, taken in 2 separate exposures).

The quality of the product data is not monitored individually (only saturation is flagged) but collectively as a process. The most important figure of merit is the achieved signal-to-noise ratio (SNR) as function of the collected photons. In Figure 2 we display the achieved SNR versus the mean of the reduced spectrum before flux calibration (after extraction and flat-fielding), for a selected setup. The mean flux is proportional to the raw signal.

We see that depending on the collected (raw) signal, the achieved SNR is following the square-root law expected if the noise is photon-noise dominated. This gives us confidence that the reduction process is indeed exploiting the full scientific information contained in the raw data. Any mismatch of calibration data and science data, for instance induced by a misalignment of the grating, would result in a loss of SNR.

For the flux calibration, we have decided to use master response curves which have been constructed from pipeline products of many individual flux standard stars. Once constructed, they have been used continuously until a major intervention (like a detector replacement) called for a new set. In that way we cannot establish an absolute flux scale since both transparency variations during the night and slit losses get ignored. But the UVES spectrograph is not aiming at spectro-photometric accuracy, and therefore the relative flux calibration seems appropriate.

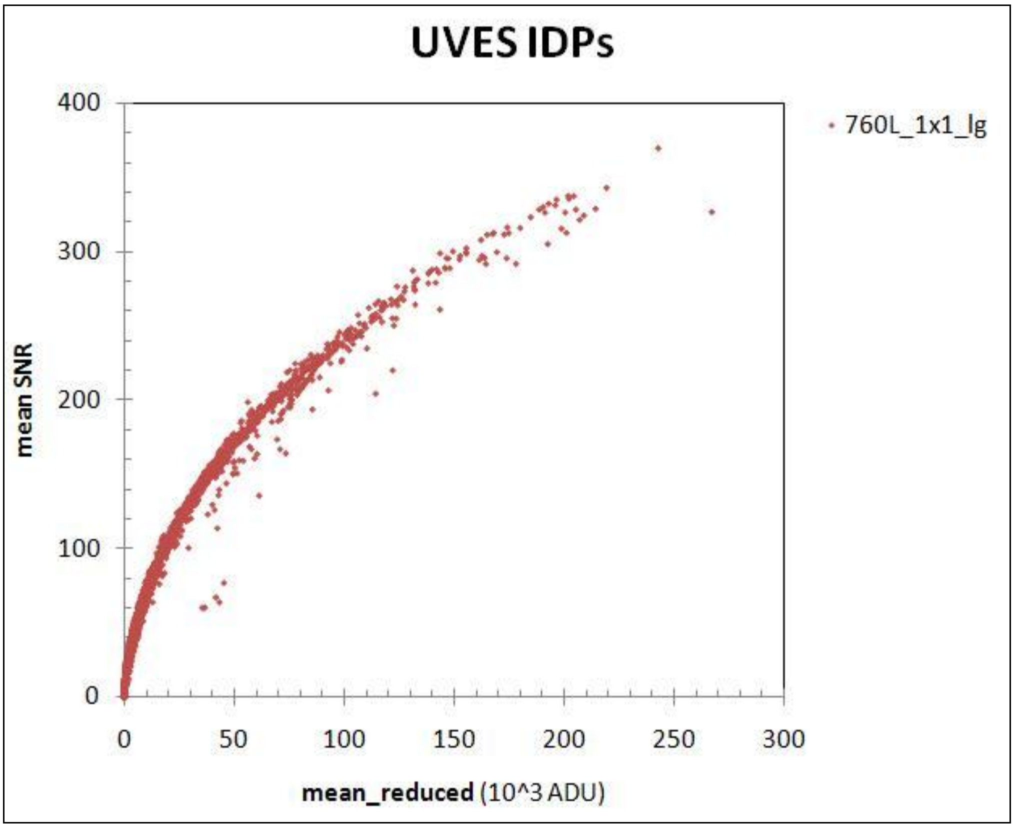

Figure 2. Phoenix process control. The QC parameters mean_SNR and mean_reduced (both averaged across the whole spectrum) are plotted here for the whole UVES date range (2000-2014), selected for the red lower detector, setting 7600 Å, 1x1 binning, low gain read mode. A total of N=7,248 data points is plotted here.

Since the opening of the archive for UVES IDPs in October 2013, we have seen more than 200 individual download requests for a total of 18,500 UVES files. The average download request is for about 100 UVES files, and there is almost one download request per day. On top of these "normal" requests there were so far two "power" requests asking for a total of 72,000 UVES files. These requests taken altogether, almost the whole archive of UVES IDPs has already been downloaded once in just 6 months. This demonstrates the huge interest in these science-grade data products and proves the validity of the assumption that there is a community interest for homogeneously reduced data from a broad range of science use cases which is complementary to the so far existing survey data products.

The UVES IDPs continue to be processed as a data stream that follows the acquisition of the parent raw data at the VLT, and the quality checks and certification of the master calibrations, with a time delay of about one month. Once released, the data products are accessible for the PIs immediately, and for all other users after the end of the proprietary period which is one year in most cases.

6. X-SHOOTER

Inspired by the huge success of the UVES IDPs, we have decided to use the *phoenix* tool to process the data from the second very popular ESO spectrograph on Paranal, XSHOOTER. That spectrograph covers an even broader spectral range, from 320nm to 2.5µm in three arms (UVB, VIS, NIR) which are exposed simultaneously and leave no gaps. XSHOOTER employs two optical detectors and one NIR array and offers a resolving power of up to 20,000. It is operational since October 2009. It offers two main modes, SLIT (single-slit Echelle spectroscopy, similar to UVES ECHELLE), and IFU. We have processed the SLIT data and offer them under the label XSHOOTER_ECHELLE.

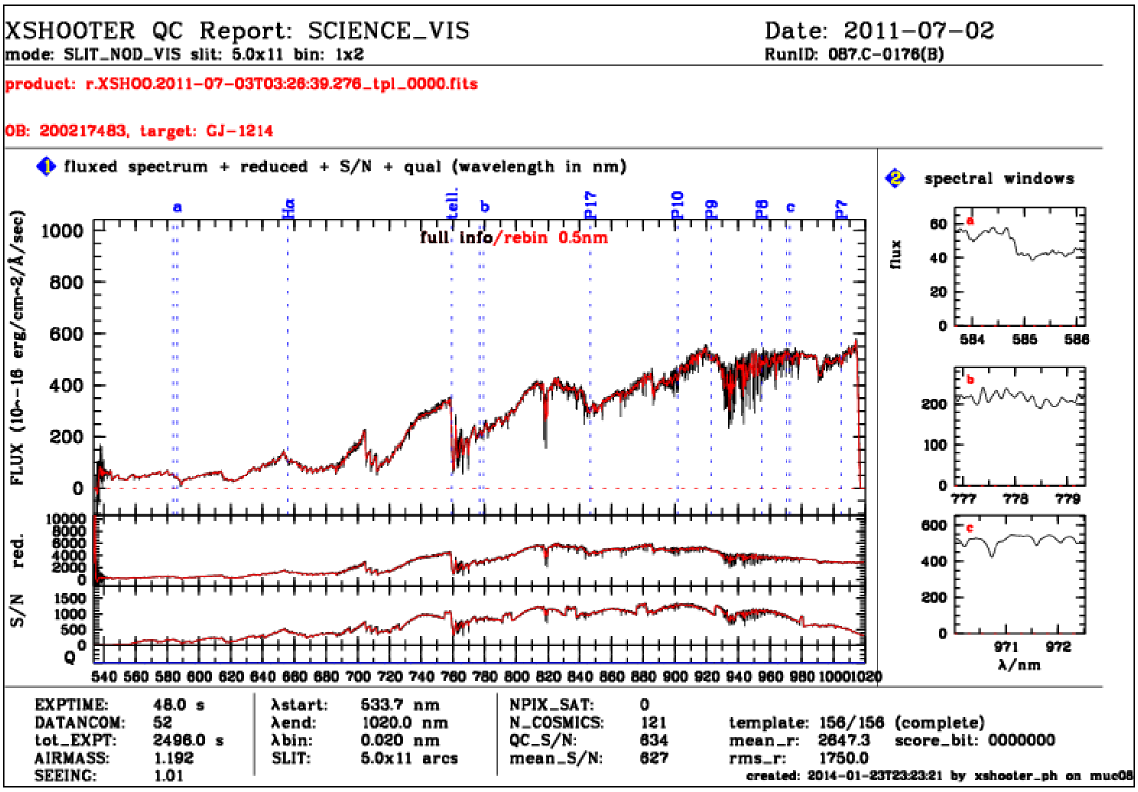

Figure 3. Example of XSHOOTER IDP. Main preview plot, featuring: 1. from top to bottom the flux-calibrated IDP (both at original resolution and as smoothed version); the unfluxed spectrum; the SNR. 2. Right column: small spectral windows at full resolution. 3. Bottom panel: a set of related QC parameters.

Since May 2014, more than 33,000 XSHOOTER IDPs have been generated and archived. They can be downloaded from the same archive interface as the UVES IDPs[8]. We have also an extensive release description[9]. They come as one product per arm, in a similar 1D binary table format as the UVES IDPs. Each XSHOOTER IDPs has also some attached files, called ancillary files:

- a 2D fits file with the reduced spectrum before 1D extraction, which is useful for quality assessment;

- a preview plot which can also be used for quality assessment;

- and a text file with additional information about the observational circumstances (observing grade and comments as given by the observatory nighttime staff).

The XSHOOTER IDPs have been processed with a reviewed and certified pipeline which saw significant improvements of the code. The data have been processed to the highest-possible science-grade level including the flux calibration with master response curves. The only major step missing is the correction for telluric absorption for which no automatic solution could be found.

The XSHOOTER instrument offers some more complexity for the observing strategy, with three possible modes: STARE, NODDING, OFFSET. There are certain assumptions associated with these different strategies: in NODDING there should always be an even number of nodding positions; in OFFSET mode there should always be pairs of SKY and OBJECT files, etc. Under the prevailing conditions during the night it was not always possible to comply with these assumptions: templates needed to be aborted or interrupted. Also, some observers tend to use the possible configurations in a creative sense. This altogether makes up for a data processing situation which is more challenging and less controlled than in the case of UVES. It was therefore unavoidable that some XSHOOTER IDPs have essentially zero flux (which could mean a failure of the extraction or e.g. the lack of signal in the raw data) or even negative flux (when sky and object positions got confused during the observation).

These data products therefore need some more care when used for scientific analysis. The assessment by the user is assisted by the delivered ancillary files, among them a quality plot (Figure 3). There is also a set of quality flags delivered in the header of the products, assigning scores 0 (for compliance) or 1 to seven key parameters, among them flux preservation, seeing-slit width comparison, saturation alert etc.

A second quality plot is delivered for the middle (VIS) spectrum (Figure 4), displaying the total spectrum from all three arms (heavily smoothed), in order to get a quick impression about the flux calibration and the overall spectral slope. It also displays the cross-dispersion profile of the signal, derived from the 2D file.

Following the example of the UVES IDPs, we have stored QC parameters in a database in order to monitor the health of the reduction process[10]. We have plotted in Figure 5 a SNR plot similar to the one for UVES. Naturally, the scatter for the XSHOOTER data is stronger, due to their larger spectral range. Nevertheless, we can demonstrate the process health with a simple trick: by selecting data points from individual programmes we can hope to collect targets with comparable spectral slope. Indeed, as shown in Figure 6, the resulting slopes for the selected three cases nicely follow the expected correlation, which is reassuring for the health of the XSHOOTER phoenix process.

---

[8] http://archive.eso.org/wdb/wdb/adp/phase3_spectral/form . Select 'any' for Programme, and 'XSHOOTER_ECHELLE' for collection.

[9] Release description: http://www.eso.org/observing/dfo/quality/PHOENIX/XSHOOTER/processing.html

[10] Browsing interface: http://archive.eso.org/qc1/qc1_cgi?action=qc1_browse_table&table=xshooter_science_public. Plotter interface: http://archive.eso.org/qc1/qc1_cgi?action=qc1_plot_table&table=xshooter_science_public

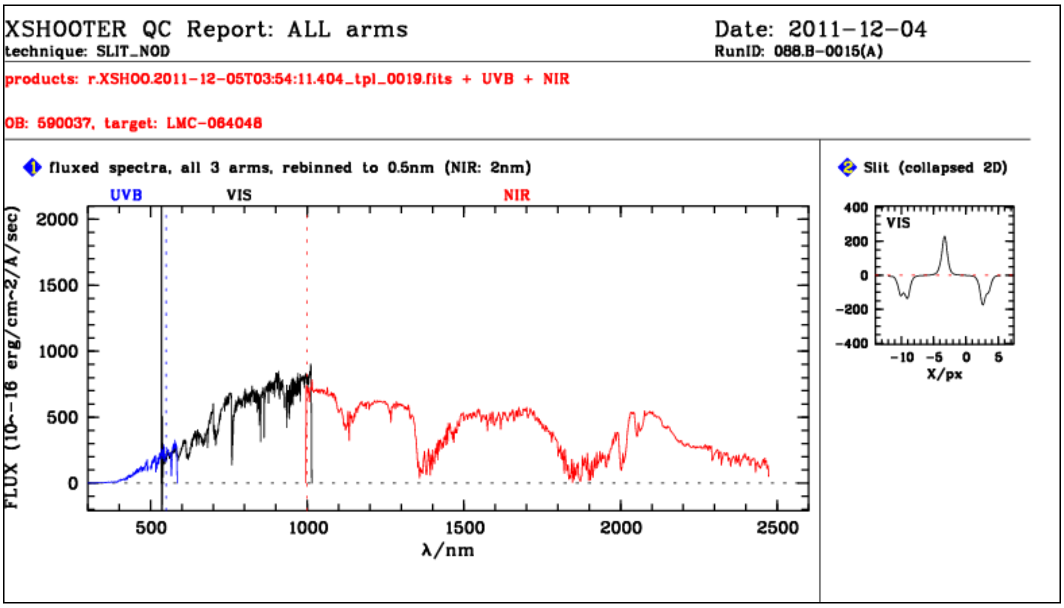

Figure 4. Second overview plot, featuring the smoothed spectral signal in all three arms (1), and the collapsed signal in cross-dispersion direction, as derived from the 2D product (2). This plot can be used to check the flux scale (effect of slit losses) and the cross-dispersion profile.

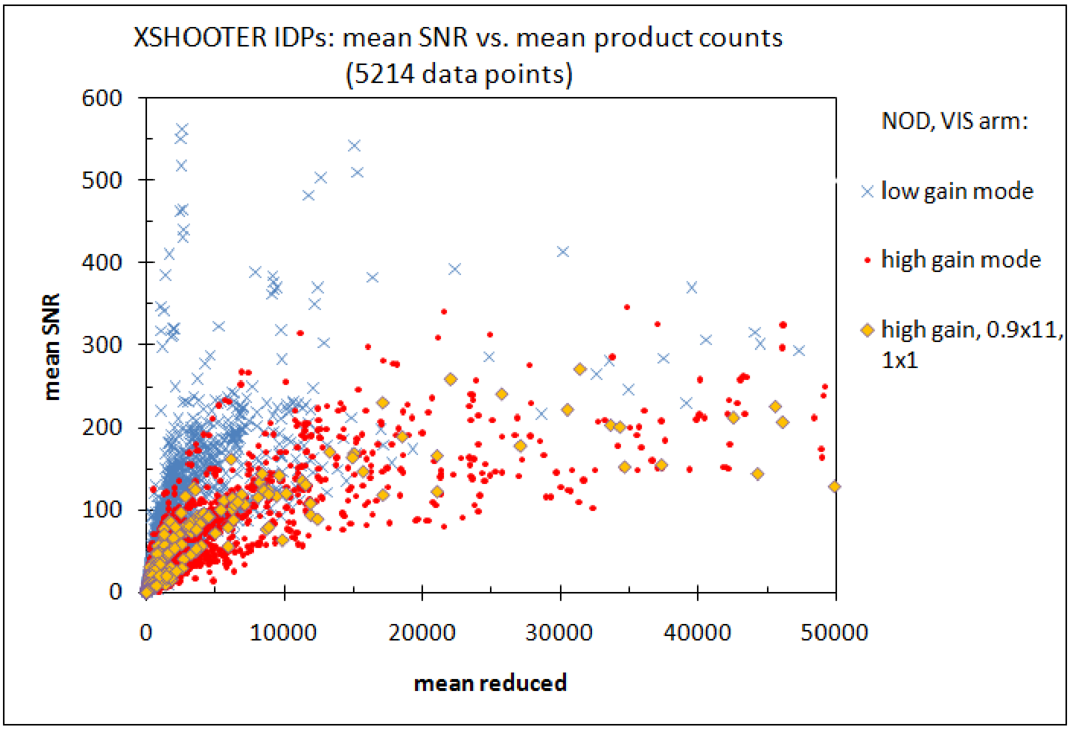

Figure 5. XSHOOTER process health control. Mean SNR vs. mean reduced spectrum, for all NODDING products of the VIS arm (a total of 5214 data points corresponding to spectra taken between 2009-10-01 and 2013-12-31). We distinguish between data for the high-gain mode (the standard mode), the low-gain mode, and a reference setting, high-gain mode, 1x1 bin, 0.9x11 slit (the most frequently used setting). The *phoenix* process for low gain mode data is able to deliver a SNR of almost 600. The scatter for the data points is stronger than for the UVES case because both parameters are averaged across a whole spectrum, and XSHOOTER spectra cover a wider range than UVES spectra. Therefore, the averaging also includes an average of the spectral slope which is of course very different among the product spectra.

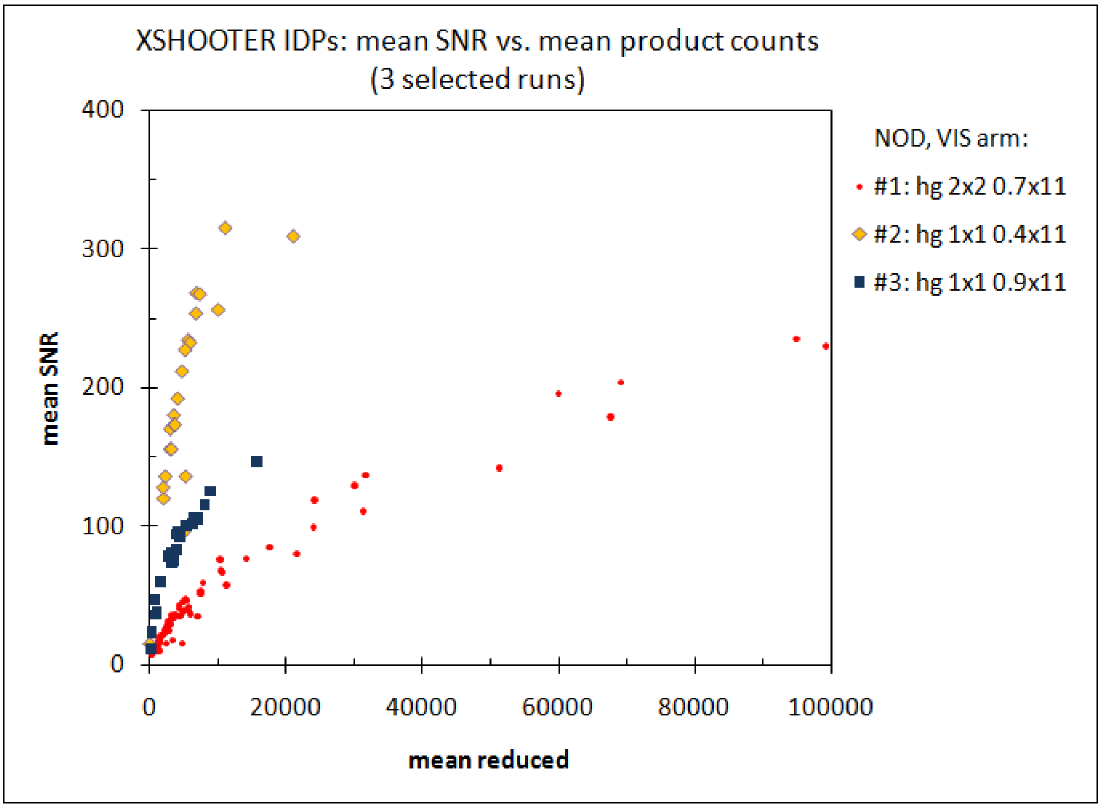

Figure 6. Same as Figure 5, for three selected runs (again NODDING products from the VIS arm only). Runs #2 and #3 have delivered data for young stellar objects, while run #1 belongs to panel D (stellar evolution). Selecting targets of similar spectral slope, we find the same well-constrained correlation between collected signal and SNR as for the UVES IDPs.

7. NEXT STEPS

Given the success of the UVES spectral IDPs, and after having finished the second major *phoenix* project, XSHOOTER, we will likely proceed to the multi-object mode of FLAMES-GIRAFFE, the MEDUSA spectra. These are up to 120 individual medium-resolution spectra obtained in one shot with the FLAMES spectrograph at the VLT. We will then extend the phoenix coverage to IFU data, as soon as a spectroscopic data standard for ESO phase 3 IFU data is available. This will bring the second-generation instruments KMOS and MUSE on the IDP path, meaning their processed data will become available through the ESO archive. Another evolutionary path for phoenix will be the second-generation VLTI instruments like GRAVITY and MATISSE.

Acknowledgements. The UVES and XSHOOTER IDP projects at ESO are the outcome of the efforts of many people at ESO, among them: Andrea Modigliani, Artur Szostak, Sabine Moehler, Wolfgang Hummel, Alberto Micol, Joerg Retzlaff, Vincenzo Forchi and Alexis Huxley.

REFERENCES

[1] Hanuschik R., Neeser M., Hummel W., Wolff B., “Scoring – a novel approach toward automatic and reliable certification of pipeline products”, SPIE 7016, [70160Q] (2008).
[2] Retzlaff J., et al., “Data products in the ESO science archive facility”, these proceedings [9149-2] (2014).
[3] Arnaboldi M., et al., “The ESO Public Surveys – Review of milestones and completion”, ESO Messenger 154, 18 (2014).