stead at the intrinsic planetary emission in the K and eventually L band.

(e) First Generation Upgrade

In addition, we have received from Jean-François Donati (Observatoire Midi-Pyrénées) a VLT instrument upgrade proposal, aimed at turning UVES into a high-resolution spectro-polarimeter. The Consortium also includes the Landessternwarte Heidelberg, Observatoire de Paris, University of St Andrews and University College London.

III. Next Steps

Following recommendations at the last April STC meeting, we are currently contacting the Project Principal Investigators. The ESO staff responsible for the four domains are respectively A. Moorwood for the KMOS, G. Monnet for the Surveyor, S. D'Odorico for the Fast Shooter (& UVES upgrade), and N. Hubin for the Planet Finder. Following current negotiations with the Consortia, feasibility studies will start soon, with (partial) financing from ESO. On a longer scale, the goal is obviously to get efficient Teams to conduct these huge endeavors.

On the programmatic side, the intention is to initiate in the next 8-year period the development of one second generation instrument per year, starting in 2003. This is of course a rough guide only. The actual rate will depend on the capital and human resources cost of the selected facilities. It is thus quite appealing that a number of Consortia intend to raise a significant fraction of the capital cost for these expensive facilities.

It may be nevertheless worthwhile to recall that a huge effort still remains to be invested to complete and put into operation the remaining first generation VLT instruments; this is heavily taxing the present capabilities of both ESO and its community for "fast" deployment of these new, exciting, facilities.

In Conclusion, I would like to again extend our deep thanks to all proponents for their valuable contributions.

VLT Quality Control and Trending Services

R. HANUSCHIK, ESO (UVES QC Scientist) and
D. SILVA, ESO (Head Data Flow Operations Group)

Why control quality?

Any observatory worldwide has staff who permanently look into the performance of the instruments and check the quality of the data. ESO's Very Large Telescope has the operational model of a data product facility. This means that it goes beyond the day-to-day performance checks and promises to deliver data of a defined and certified quality.

The Data Flow Operations Group in Garching (DFO, also frequently called QC Garching), provides many aspects of data management and quality control of the VLT data stream. One of the main responsibilities is to assess and control the quality of the calibration data taken, with the goal to know and control the performance of the VLT instruments. Information about the results of this process is fed back to Paranal Science Operations and to the ESO User Community via QC reports and web pages.

The constant flow of raw data from the VLT instruments splits into data streams for the science data and the calibration data. The calibration data stream has two separate components:
- calibrations taken to remove instrument signatures from science data
- calibrations taken for routine daily instrument health checks.

The focus of QC Garching is to process these data and extract Quality Control information. This process of course does not replace the on-site expertise of the Paranal staff. But it goes beyond the usual quick-look, on-the-spot checks and provides a permanent and in-depth knowledge of the instrument status. With the QC parameters routinely collected over years, it is possible to control, predict, and often even improve, the performance of the instruments.

This article describes the Quality Control process for the four presently operational VLT instruments: FORS1+2, ISAAC and UVES. This process will be extended and refined for the next suite of instruments coming soon, VIMOS, NACO, and FLAMES, and ultimately expanded to all VLT instruments.

How to control data quality

The term quality control, though often used, needs some definition. Quality control, as we understand it, implies the control of the following:
- quality of the raw data
- quality of the products and of the product creation process
- performance of the instrument component involved.

Quality control does generally not imply aspects like the monitoring of ambient data (quality of a night), the proper format of FITS headers, or the tracking of programme execution. Responsible for these aspects, being part of Quality Control in a wider sense, are other groups, e.g. Paranal Science Operations, and the User Support Group.

Pipelines. Fundamental in the QC process is the use of automatic data processing packages, the pipelines. Without these, effective quality control of the huge amount of data produced by the Observatory would be impossible. In fact, the primary goal of the data reduction pipelines is to create calibration products and support quality control. Only after this comes the reduction of science data.

With the large-scale use of data processing pipelines, the Quality Control group has effectively also the function of assessing and improving the accuracy of the pipelines. As a by-product, we provide documentation about the pipeline functions from the user's point of view.

The usual day-to-day workflow of the QC scientists has as primary components: processing the raw data (calibration and science) using the instrument pipeline, performing the quality checks, and selecting the certified products and distributing them.

One might say that the use of the pipeline, once the process has been set up properly, is mainly number-crunching, while the quality checks require expertise.

The QC process. There is a natural three-floor pyramid in the QC process (Figure 1):

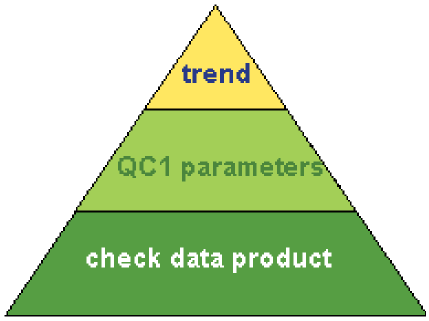

Figure 1. The QC and trending pyramid.

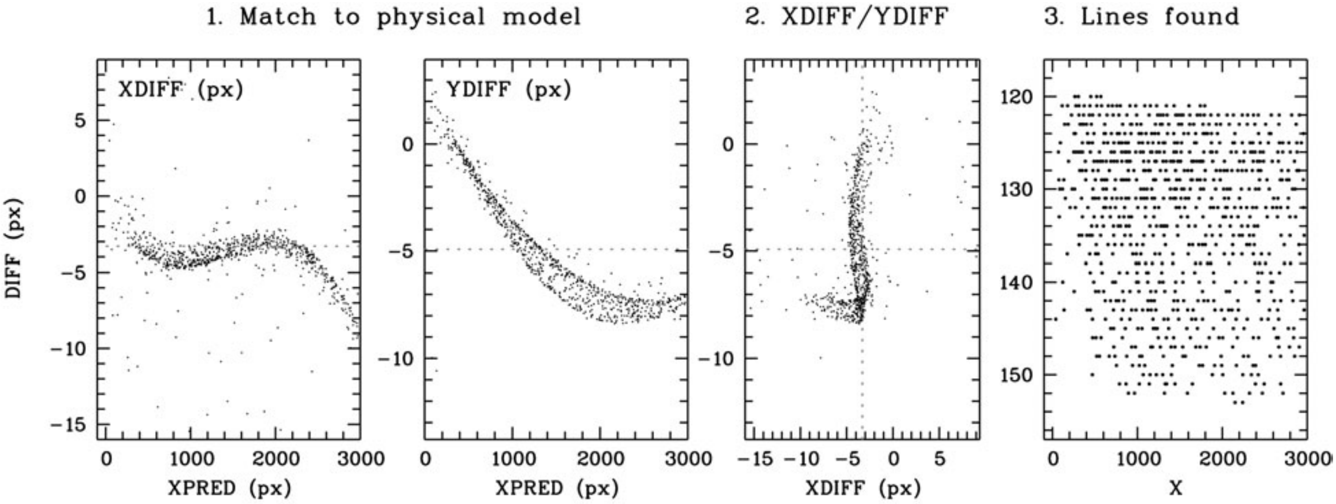

Figure 2. Quality plot for a UVES "formatcheck" frame. Such frames are taken daily to control the proper adjustment of the gratings and cross-dispersers. Main focus here is the proper clustering of the line positions found (boxes 1 and 2 with the difference between predicted and found line positions) and the proper coverage of all orders with identified positions (box 3).

- checking each data product
- measuring and storing a set of QC parameters per product, and
- looking at the long-term behaviour of these parameters.

1. Checking each product. The first, and most fundamental, step in the QC process is done on each pipeline product. Is the frame over/under-exposed? Is it different from the frame taken yesterday? A set of procedures creates displays and graphical information like cross-cuts and histograms. Frequently there is a comparison to a reference frame. Without these procedures, the QC scientist would be blind to data quality. The mere fact that a reduction job was executed without a processing error indicates nothing about data quality.

In practice, after some initial phase when indeed everything is inspected, one usually decides to switch to a 'confidence mode' in which, for example, only every third night is inspected in depth, while for the others the trending plots are consulted. This strategy is economic in case of very stable instrument performance, and mandatory with high data rates. Then the 'human factor', namely the possible level of concentration, ultimately limits complete product checks.

Figure 2 shows as an example the QC plot for the products of a UVES *formatcheck* frame which is a technical calibration needed by the pipeline to find the spectral format. With an experienced eye, just a second is needed to know from this plot that everything is fine and under control.

QC checks are also done on Paranal, directly after frame acquisition. These on-the-spot inspections are of quick-look character and apply to both raw and product data. They are extremely important to check the actual status of the instrument, especially for those instruments like UVES which have the data transferred to Garching via airmail.

2. Deriving QC1 parameters. The next step in the QC process is the extraction of QC parameters. These are numbers which describe the most relevant properties of the data product in a condensed form. Since they are in most cases derived through some data manipulation (e.g. by the reduction pipeline), they are called QC1 parameters. This distinguishes them from the QC0 parameters which mainly describe site and ambient properties like seeing, moon phase etc.

Across the instruments, there are always QC1 parameters describing the detector status, i.e. the read noise, the mean bias level, the *rms* of gain variations etc. Other QC1 parameters specific to spectroscopic modes are resolving power, dispersion *rms*, or number of identified lines. Imaging modes are controlled by QC1 parameters like zeropoints, lamp efficiency, and image quality.

3. Trending. The top level in the QC pyramid is the trending. Trending is a compilation of QC1 parameters over time, or a correlation of one QC1 parameter against another one. Trending can typically prove that a certain instrument property is stable and working as specified. It can do much more, however. For example, trending can discover the slow degrading of a filter, or aging effects of the detector electronics.

Outliers. Within the trending process, main attention focuses on two extremes: the outliers, and the average data points. Information theory says that outliers transport the highest information content. But not all outliers are relevant for QC purposes. We need to distinguish whether the outlier comes from a bad algorithm setup, from a bad instrument setup, or just from a bad operational setup.

A bad algorithm may, for example, be a wrong code for *rms* determination. Such an outlier would help to improve the code. A bad instrument setup could be a stuck filter wheel with the filter vignetting the light path. A bad operational setup would be a frame labelled as a *flat*, but with the lamp not switched on. In short, evaluation of the trending data is non-trivial and requires judgement.

Finding that a certain QC1 value is stable over months or years may lead one to relax the acquisition rate of the corresponding calibration data. This may be a good idea since we should avoid over-calibration. But one has to bear in mind that for proving stability, one needs a good coverage in time, so it is a good idea to have calibrations done more frequently than their typical variation timescale.

Certification. Once a product file has been QC checked, and its QC1 parameters have been verified to be valid entries, the data enter the delivery channel, which involves ingestion into the master calibration archive, usage for science reduction and distribution to the end users (if taken in Service Mode). By definition, the data are then certified. Rejected data are deleted.

Components controlled

UVES has been operational since April 2000 and was the first instrument with built-in QC procedures as part of the pipeline. So it was possible to measure and collect from the beginning a backbone set of QC1 parameters. This set has been expanded over the last two years and is now almost complete. The two-year baseline forms an asset from which many valuable pieces of information about long-term behaviour can be extracted.

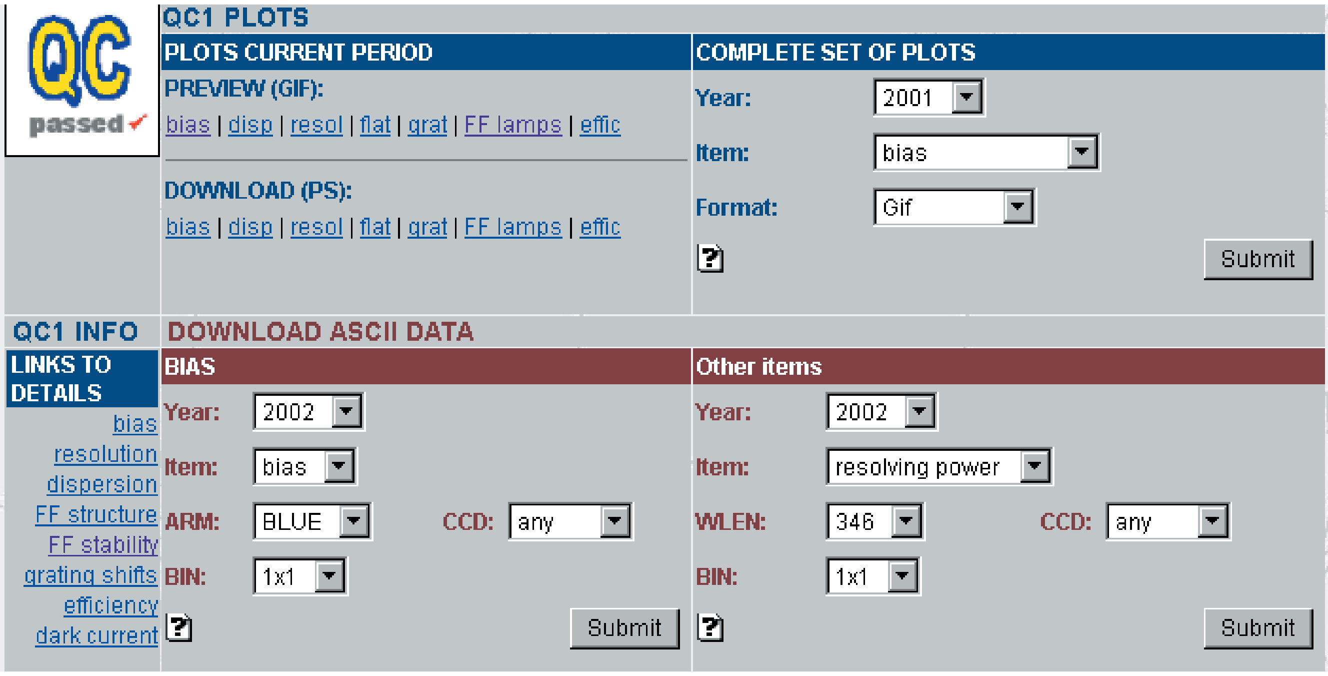

*Figure 3. This web interface (*http://www.eso.org/qc/UVES/qc/qc1.html*) connects to the QC1 data of UVES. The user may view trending plots, and download the corresponding ASCII data. A quick-look panel for the current period links to all current trending plots, i.e. those which are relevant for the present instrument health.*

Trending

With the QC1 parameters stored in a database they become available for trending. There is a central QC1 database under development which will host all QC1 parameters and other related information such as plots and trending results.

The QC1 database can be considered as the central memory about the status of each VLT instrument. The goal is to have available all quality information from the complete operational history of the instrument. This also includes information about interventions (e.g. mirror recoating) and replacements (optical components, detectors). Such information is vital for interpretation of trending results. Moreover, with data collected over years, it becomes possible to detect slow degrading effects. Preventive interventions and maintenance can be scheduled properly.

Web pages

As part of the QC process, these results are published on the web. The central Quality Control site is http://www.eso.org/qc/ which has, per instrument, a link to QC and trending results.

Under the URLs http://www.eso.org/qc/<INSTR>/qc/qc1.html (where INSTR is any of UVES, FORS1. FORS2 or ISAAC), one connects to the QC1 database and views trending plots and ASCII data (Fig. 3). Here one also finds detailed documentation about the QC1 parameters.

Our goal is to present knowledge, not just information. Take as an example the trending of the UVES spectral resolving power *R*. We do not just dump all available numbers per date, but provide a documentation of the measurement process, a selection of trending plots, a correlation with slit width and a comparison to User Manual values.

UVES QC monitors the following instrument components (http://www.eso.org/qc/UVES/qc/qc1.html):

Component	Property
detector	bias level, read noise, dark current; fringing
gratings	stability of spectral format; resolving power, precision of dispersion solution
slit	slit noise
lamps, filters	FF lamp stability, filter throughput
all components	efficiency

FORS1/2 have the following QC items (http://www.eso.org/qc/FORS1/qc/qc1.html, the FORS2 link has replaced “FORS1” by “FORS2”); MOS will be added eventually:

Mode	Property
General	bias level, read noise, dark current, gain, contamination
Imaging	zeropoints, colour terms, image quality
Long-Slit Spectroscopy	dispersion, resolution

ISAAC has the following QC items (http://www.eso.org/qc/ISAAC/qc/qc1.html), with the long-wavelength arm being added soon:

Mode	Property
General	dark level, read noise
Imaging (short arm)	Zeropoints

Some QC and trending highlights

Compensation of UVES thermal drifts. The precision of the UVES spectrograph is limited by ambient temperature changes. A one degree difference causes an effective shift of the gratings by up to 1 pixel in cross-dispersion direction (Figure 4). The daily *formatcheck* frames are compared to a reference frame and used to measure these shifts. Since the QC1 values proved a general stability and a linear slope of the thermal coefficients, a compensation for such drifts was successfully implemented in cross-dispersion direction, and the dispersion direction was corrected.

UVES filter degradation. The monitoring of the exposure level of the UVES flat-field lamps gives control over the lamp and filter status. The filter status is especially significant for the quality of science observations. In July 2001, the transmission of the blue CuSO4 filter dropped, which was discovered in the trending plot (Figure 5). The replacement of the filter gave the blue efficiency of UVES a boost.

FORS1 image quality. Figure 6 combines input data from pipeline-processed science images from FORS1. It demonstrates that in most cases FORS1 image quality is determined by the seeing and not degraded by potential errors such as telescope guiding.

FORS1 zeropoints. Figure 7 shows the complete history of FORS1 zeropoints in the V band, spanning three years. Zeropoints measure the efficiency of the overall system instrument plus telescope. There have been major interventions (see the caption for details), but perhaps more interesting is the fact that there has been a general loss of efficiency by about 8% per year, due to degrading of the mirror surface.

ISAAC photometric zeropoints. Figure 8 shows the zeropoints derived by the ISAAC pipeline over a half-year period. The sharp rise around MJD-OBS = 52,200 is due to an intervention which included a re-alignment. This improved the instrument efficiency by up to 0.2 mags. The long-term trend is due to degradation, while the short-term scatter in most cases is due to fluctuations of the night quality.

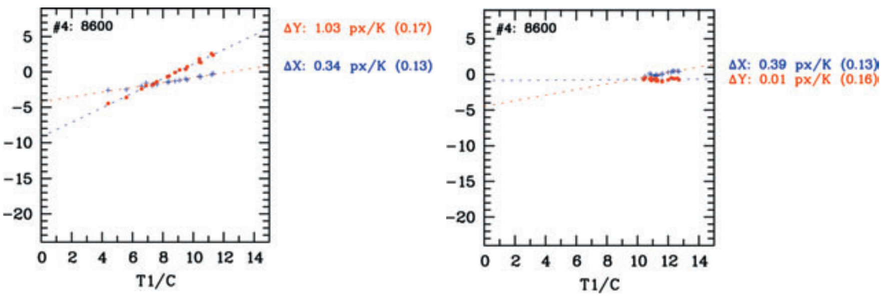

Figure 4. Measured thermally induced drifts of UVES grating #4, without (left), and with (right) thermal compensation in Y direction. The QC1 trending data have been used to establish the coefficients for the automatic compensation of thermal motion in Y (cross-dispersion) direction.

Shared QC

It is obvious that Quality Control must be a shared responsibility between QC Garching and Paranal Science Operations. There are always QC issues which require immediate reaction and intervention. These can only be properly handled on site. With the data airmailed to Garching (which today is the transfer mode for UVES, and soon for all VLT instruments), the typical reaction time on QC issues in Garching is 10 days. This naturally leads to the concept of shared QC which means that part of the QC tasks are done on Paranal (in real time, by daytime astronomer), part in Garching (off-line, by QC scientist).

On-site QC. Basic quality checks on the calibration data are performed by the Paranal daytime astronomer. Just after exposing the raw calibration data and pipeline-processing them into calibration products, the data are inspected visually. The on-line pipelines produce an essential set of QC1 parameters which is fed into a database. Essential are those QC1 parameters relating to fundamental instrument properties which, in case of failure, would jeopardize the usefulness of the science data. Such instrument health parameters are e.g. proper adjustment of gratings and filters, and proper CCD setup.

Off-line QC. The full set of quality checks is applied in Garching – anything which is not time critical, but requires in-depth analysis, pipeline or post-pipeline procedures. This applies also to complex trending analyses requiring extended data sets. Examples are photometric zeropoints which are determined from all standard star data of a night; colour and extinction terms being derived for a whole semester; efficiency curves; sky brightness etc.

Feedback loops. The exchange of quality information between the two sites with QC activities is especially important. The main feedback channel from QC Garching to Paranal are the web-published trending plots which are updated daily. These monitor the proper function of all QC-checked components. Any anomalies are investigated in detail and reported directly to Paranal.

IOT. The group responsible for studying new results of trending studies, the development of new QC1 parameters or algorithms are the Instrument Operating Teams (IOT). As a vertical structure, these have, per instrument, delegates from each team which is essential for proper operations, i.e. from Science Operations, User Support Group, DFO, Pipeline Development, and Instrumentation. The teams are led by the Instrument Scientist. It is here where all the expertise about a VLT instrument is focused and where the loop between QC results and improvement of instrument performance is closed.

Feedback into calibration plan

An example of successful feedback of the QC process and the instrument operations is the improvement of the Calibration Plan. In principle, calibration data are taken to remove instrumental signature from the science data ("calibrate the science"). In a more general sense, they are taken to know the instrument status ("calibrate the instrument"). Ideally, one would go for the latter goal since this provides the broader knowledge about the instrument while including the requirement to reduce the science data.

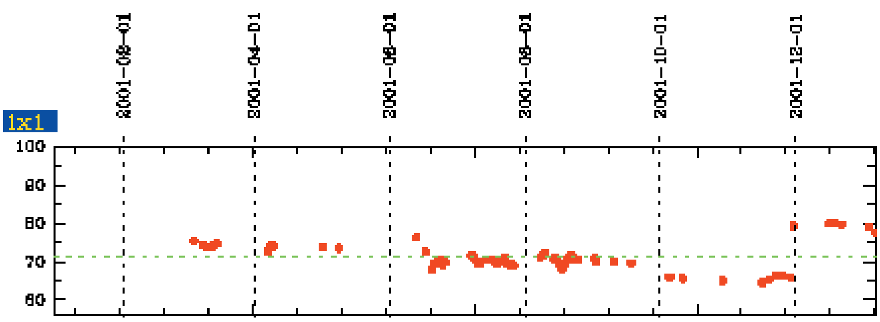

Figure 5. The transmission of the CuSO4 filter used for reducing scattered light in the blue arm of UVES dropped significantly in July 2001. This was only discovered in November 2001 when the corresponding trending procedure had been established. An inspection of that filter verified its poor state: its coating was partly destroyed by humidity. Its replacement in December 2001 has improved the efficiency considerably, which is clearly visible in the trending plot.

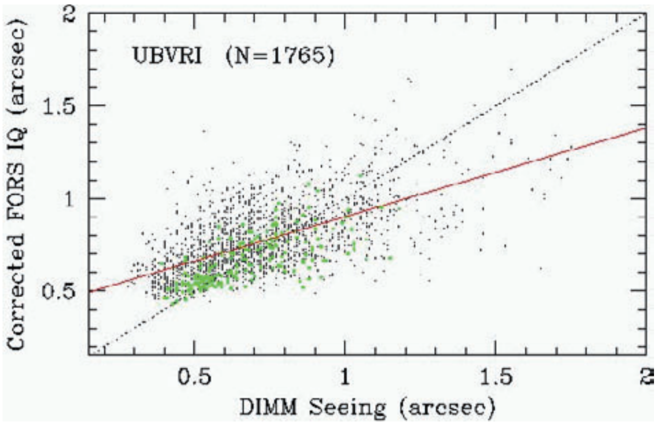

Figure 6. Image quality of FORS1 (width of stellar images in arcsec) versus the image quality measured by a seeing monitor (DIMM). Input data have been collected from pipeline-processed FORS1 science images in filters UBVRI. Correction factors have been applied for wavelength and airmass. Green dots mark high-resolution collimator data, black stands for standard-resolution. The broken line indicates FORS IQ = DIMM. The red line is a fit to the data. FORS1 image quality is on average better than DIMM seeing above 0".8.

But in practice this is not even possible for simple instrument modes. For instance, the imaging mode of FORS1 has 5 standard filters with 4 CCD read modes and 2 different collimators. Obtaining a complete set of calibration frames, including twilight flats and standard stars, is practically not possible every night. As a more complex instrument, UVES has 12 standard setups, with roughly 20 different slit widths and 2 CCD modes, and the parameter space becomes forbiddingly large for routinely calibrating all settings.

Therefore, calibration data are usually triggered by the science setups actually used in the night before. To these are added the daily health check calibrations. On Paranal, an automatic tool is used which collects this information into the daytime calibration queue.

Even this strategy still produces a large calibration overhead both in terms of exposure time and archive disk space. So, after some initial epoch when confidence has been gained that the calibration plan is complete, one may start thinking how to *optimize* the plan. In the case of UVES, it has been shown by the trending studies that the most relevant instrument properties usually show trending timescales much longer than a few days. Based on this experience a three-day memory has been implemented in the calibration plan, with calibrations for an identical setup being repeated only every three days. The only exception is the wavelength calibration. The health check calibrations are executed daily in order to prove that nothing irregular happens, e.g. an earthquake which would clearly break the long-term trending assumption.

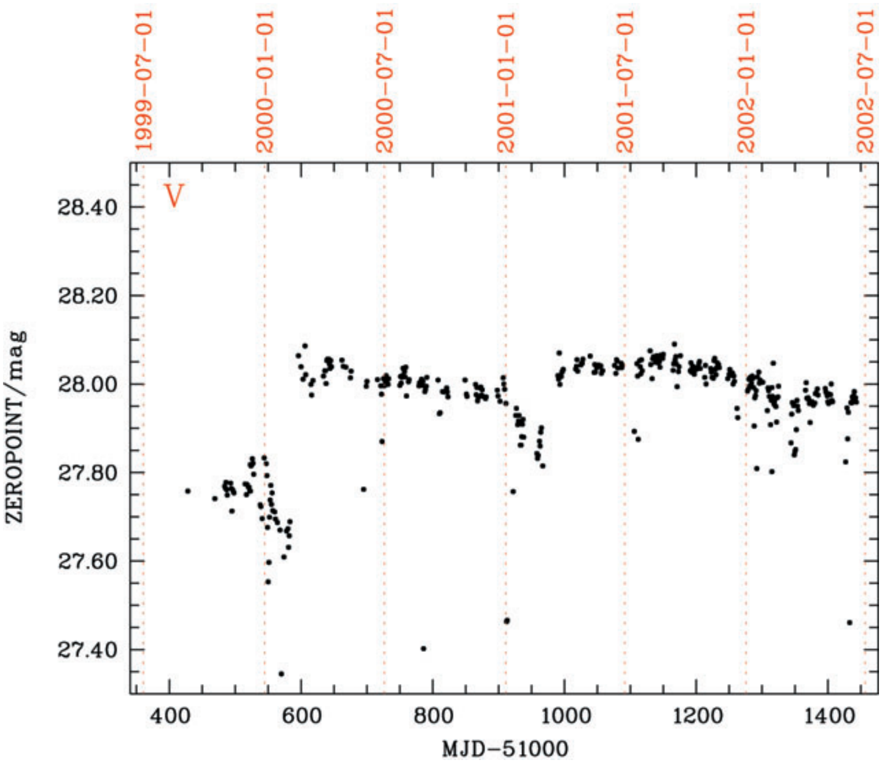

Figure 7. Three years of FORS1 zeropoints in the V band. Major interventions, causing steps in the slope, have been mirror recoatings in February 2000 and March 2001, and sudden degrading of the main mirror due to rain in February 2001. The move to UT3/Melipal in August 2001 is invisible in this plot.

Vision

Our QC process is continuously evolving, to meet the current and future needs of our main customers: Paranal Science Operations and the ESO user community. Here are a few examples.

Although the QC1 parameters computed and controlled by QC Garching are available via our Web pages, they are not easily associated with the calibration products available from the ESO Science Archive. By the end of this year, we hope to have a new QC1 parameter database within the Archive domain. Once this database exists, it should be possible for users to retrieve the QC1 parameters associated with the calibration products they are retrieving from the Archive. This is particularly important in the context of Virtual Observatory development.

The calibration data flowing through QC Garching contains a rich but largely unexploited reservoir of information about Cerro Paranal as a site. QC Garching, in collaboration with other group within ESO, has started several projects to process this information and make it available to our customers. For example, this year we will publish a high signal-to-noise, high resolution sky atlas extracted from many hours of UVES observations, as well as a study of optical sky brightness as a function of lunar phase, lunar distance, time after twilight, etc, derived from FORS data. Possible future projects include the creation of lists of faint, secondary photometric standards for FORS and ISAAC, in collaboration with Paranal Science Operations.

For historical reasons, our QC web pages (http://www.eso.org/qc) are implemented in a very heterogeneous way. We are reorganising and revising these pages to make them more homogeneous across instruments, and to make it easier for our users to find the information they need.

Of course, our main priority this year is to establish regular QC operations for the latest VLT instruments: NACO, VIMOS, and FLAMES, as well as extending our process to the VLT Interferometer complex. These instruments introduce many new and complex modes: optical interferometry, adaptive optic imaging, high density multi-object spectroscopy with slits and fibers, and integral field spectroscopy. The underlying, detector based health and wellness QC process are essentially extensions of our current process, but the development of a higher level QC process will be more challenging.

Last Words: Other DFO Services

In this article, we have focused on our main role: quality control for the VLT data flow. To close, we briefly outline the other services DFO provides to the ESO user community.

As mentioned above, the VLT QC revolves around calibration data. Most quantitative QC is done using calibration products, e.g. dispersion solutions or master flat fields. It is the responsibility of DFO Garching to produce such calibration products and then re-use them in a number of ways. Calibration products are ingested into the ESO Science Archive; they are included in the standard data packages produced for Service Mode users; they are used within the on-line DFS Pipeline system on Paranal; and they are used to produce science data products for Service Modes users.

Another important DFO service is processed Service Mode science data. Science data are only processed when an appropriate pipeline is available. Science pipeline data products are delivered with the understanding that they may not be publication quality in all cases. However, these products can be very useful for making the initial assessment of science data quality and for providing guidance on how to process the delivered science data for a specific application.

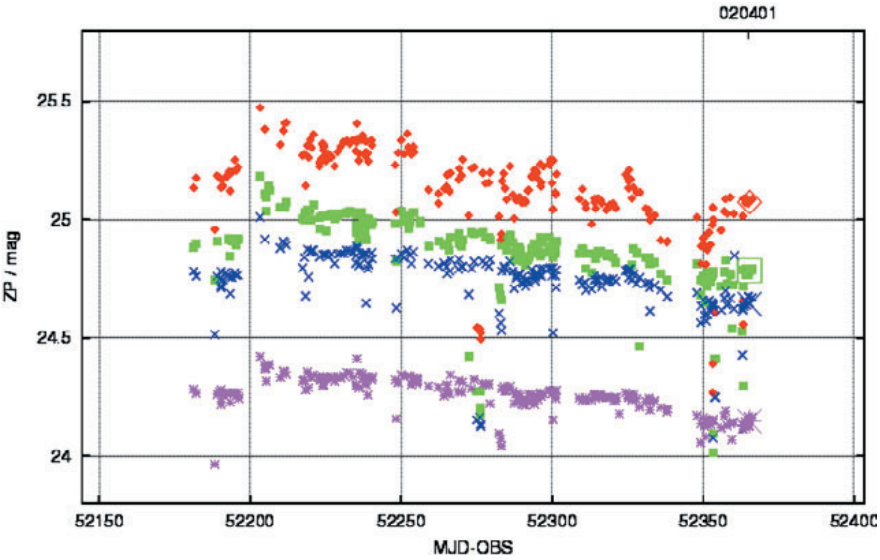

Figure 8. ISAAC photometric zeropoints for Period 68, in photometric bands Js (green squares), J (red diamonds), H (blue crosses), and K (asterisks). Horizontal axis is Julian Date of observation (MJD-OBS), vertical axis is zeropoint in magnitudes. Last civil date of observation on the plot is 2002-04-01.

When a Service Mode run is completed, DFO creates and delivers a standard data package to the run Principal Investigator. This data package contains all the raw science and calibration data, pipeline science and calibration products when available, and a variety of supporting listings and reports. Technical support (e.g. media manufacturing) is provided by the ESO Science Archive team.

Last but not least, DFO maintains extensive documentation about what we do and how we do it, on our QC Web pages: http://www.eso.org/qc/. Our detailed descriptions of how science and calibration data are processed using the current generation pipelines may be particularly interesting to users.

Acknowledgements. The QC process described here is the result of the joint work of the QC Garching team which is constituted, apart from the authors, by Wolfgang Hummel, Roberto Mignani, Paola Sartoretti, and Burkhard Wolff. We also thank our past DFO colleagues Paola Amico, Ferdinando Patat, and Bruno Leibundgut, and all our PSO colleagues, especially Andreas Kaufer.

Recent NAOS-CONICA Images

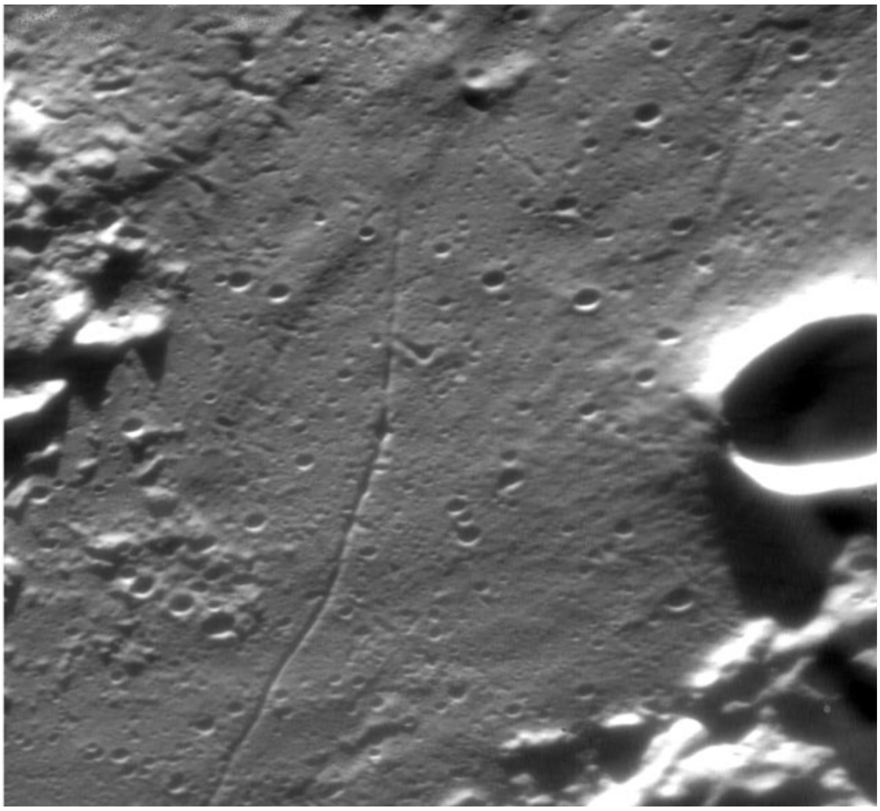

Walking on the Moon: *The ability of NAOS to do wavefront sensing on extended objects was once again demonstrated by closing the AO loop on the peak of a sunlit lunar mountain. This image was obtained with CONICA through a 2.3 micron narrow-band filter and shows details down to 0.1 arcsec (which corresponds to a ground-resolution of 175 metres, quite comparable to the resolution obtained from lunar orbit with the NIR camera aboard the Clementine spacecraft). The NACO image covers a region of about 45 × 45 km.*
(Picture credit: Eric Gendron and the NAOS and CONICA consortia.)