VISTA data flow system: pipeline processing for WFCAM and VISTA

Mike Irwin[a], Jim Lewis[a], Simon Hodgkin[a], Peter Bunclark[a], Dafydd Evans[a], Richard McMahon[a] Jim Emerson[b], Malcolm Stewart[c] and Steven Beard[c]

[a]Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, UK;
[b]Astronomy Unit, School of Mathematical Sciences, Queen Mary University of London, Mile End Road, London, E1 4NS, UK;
[c]UK Astronomy Technology Centre, Royal Observatory, Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK

ABSTRACT

The UKIRT Wide Field Camera (WFCAM) on Mauna Kea and the VISTA IR mosaic camera at ESO, Paranal, with respectively 4 Rockwell 2kx2k and 16 Raytheon 2kx2k IR arrays on 4m-class telescopes, represent an enormous leap in deep IR survey capability. With combined nightly data-rates of typically 1TB, automated pipeline processing and data management requirements are paramount. Pipeline processing of IR data is far more technically challenging than for optical data. IR detectors are inherently more unstable, while the sky emission is over 100 times brighter than most objects of interest, and varies in a complex spatial and temporal manner. In this presentation we describe the pipeline architecture being developed to deal with the IR imaging data from WFCAM and VISTA, and discuss the primary issues involved in an end-to-end system capable of: robustly removing instrument and night sky signatures; monitoring data quality and system integrity; providing astrometric and photometric calibration; and generating photon noise-limited images and astronomical catalogues. Accompanying papers by Emerson *et al.*[1] and Hambly *et al.*[2] provide an overview of the project and a detailed description of the science archive aspects.

Keywords: image processing, cataloguing, quality control, calibration, WFCAM, VISTA

1. OVERVIEW

With a typical expected average data rate of an image of the sky every 5–30s, the wide field NIR mosaic cameras on WFCAM and VISTA present both a challenge and an opportunity. The challenge is to optimise the data taking strategy and to provide an automatic end-to-end data processing and effective archive system for the several hundred Gbytes of data taken each night. The opportunity is to survey the NIR sky several magnitudes fainter than currently achieved, thereby opening up exciting scientific opportunities across a broad spectrum of contemporary astronomical research.

Although we are mainly concerned here with describing the pipeline architecture being developed to deal with the IR imaging data from WFCAM and VISTA, we would like to emphasise the integrated nature of the end-to-end system; in particular, the key role played by careful design of observing protocols and the use of associated FITS header keywords to drive the automatic pipeline processing. Figure 1 illustrates the nested nature of the data taking philosophy. We have accordingly designed the pipelines around a selectable modular scheme driven by processing recipes for maximum flexibility.

Our general philosophy is that all fundamental data products are FITS multiextension files (MEFs) with headers describing the data taking protocols in sufficient detail to trigger the appropriate pipeline processing components, and that all derived information, quality control measures, photometric and astrometric calibration and processing details, are also incorporated within the FITS headers. Generated catalogues are stored as multiextension FITS binary tables. These FITS headers thereby provide the basis for ingest into databases for archiving and databases for real time monitoring of survey progress and survey planning.

For all correspondence e-mail: mike@ast.cam.ac.uk

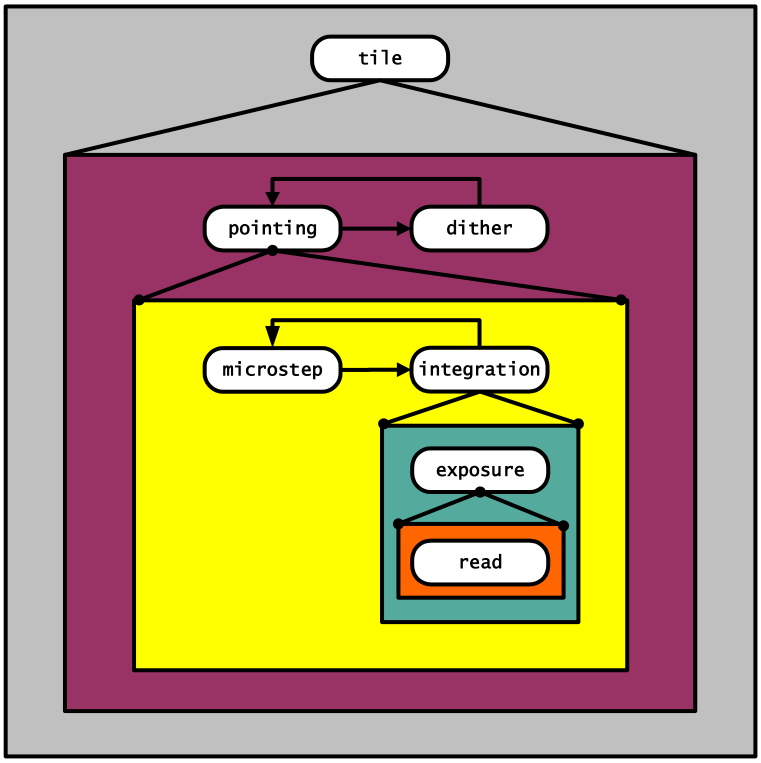

Figure 1. The underlying elements within a WFCAM observing block. Each component is an optional iteration of a sequence of observations. At the lowest level, an exposure is made up of a sequence of reads. Exposures can be looped over (*e.g.* coadded) to build up an integration. Microstepping over integrations can be used to improve sampling by interleaving the data frames. Dithering is used to remove bad pixels and improve depth. Tiling is used to fill the gaps between the mosaicked detectors to give contiguous sky converage. Offsets to other targets, or for sky estimation, are not included for clarity. Each operation is recorded in the FITS header keywords which are used to drive the pipeline.

In order to deliver against various operational and scientific requirements, there is a need for: causal summit pipelines to assess data quality on-the-fly, provide feedback to survey planning and progress tools, and provide first-pass science products for users at the summit; standard calibration and science product pipelines to be run in the UK and at ESO, Garching, to remove instrumental signatures, produce catalogues and enable astrometric and photometric calibration in addition to extra quality control checks; further processing pipelines to generate well-sampled PSFs (Point Spread Functions) and to undertake PSF and Sersic profile fitting; and an advanced database-driven pipeline toolkit for dealing with the complex issues of asynchronous image processing, *e.g.* deep stacking, merging data from different passbands, and optimal multi-passband knowledge extraction.

To reduce the data storage, I/O overheads and transport requirements, we intend to use, as much as possible, the lossless Rice tile compression scheme as used transparently, for example, in CFITSIO.[3] For this type of data (32 bit integer) the Rice compression algorithm typically gives an overall factor of 3–4 reduction in file size.

2. HARDWARE ISSUES

As an illustration, for WFCAM, summit processing will use a fast Linux PC per detector, processing data from all four detectors in parallel, and will use commodity disk RAID arrays for local storage. A 5th summit

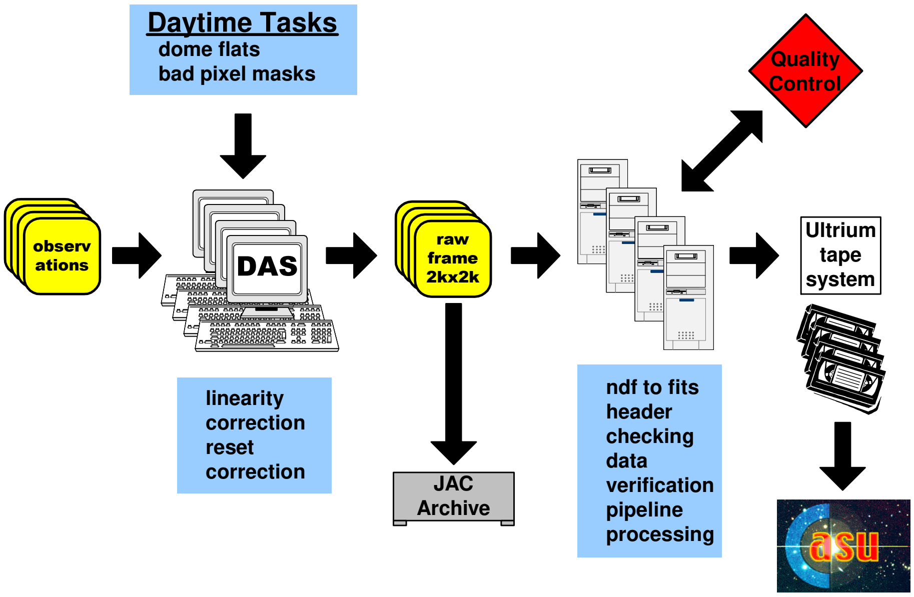

Figure 2. An overview of the summit processing setup for WFCAM. Each detector feeds a separate data aquisition PC, which then forwards the images in parallel to the summit data processing PCs. Data transport to the UK for science processing is via Ultrium LTO-I tapes.

data processing PC will be used to perform a limited set of pseudo real time astronomical analysis tasks. Figure 2 schematically shows the corresponding data flow operations at the summit on Mauna Kea.

In Cambridge, to provide a backup for the main WFCAM on-line science archive and to facilitate several of the advanced processing/reprocessing stages, we will use an on-line RAID-based storage system for the science frames and master calibration data. For WFCAM this will essentially be an extension of the rest of the existing UKIRT raw data archive currently held in Cambridge. With compression we anticipate the yearly storage requirements for WFCAM to be ≈10 Tbytes per year and for VISTA ≈50 Tbytes per year.

Transfer of processed data from Cambridge to the WFCAM and VISTA Science Archives in Edinburgh will be via the Internet. The Cambridge pipeline setup is illustrated in figure 3. Although these hardware models are directed specifically at WFCAM, they are scaleable and hence a similar solution should also suffice when VISTA comes on line in late 2006.

3. IMAGE PROCESSING

The main technical challenges in processing WFCAM and VISTA data stem from the fact that: IR detectors are currently inherently more unstable than their optical counterparts; the sky emission, roughly 100 times brighter than most objects of interest, varies in a complex spatial and temporal manner; and the large data volume that arises from NIR mosaic cameras. We highlight in the next two sections the main processing steps involved to counter these challenges.

3.1. Processing in the Data Aquisition System (DAS)

At present there are several basic reduction steps that can be done at the telescope in the DAS before data are written to disk. In the case of the first two items, these are done in order to cut the data rate by a large factor:

Reset correction: the reset image (*i.e.* bias frame) will be removed in situ, eliminating the need to write a separate image for each data frame.

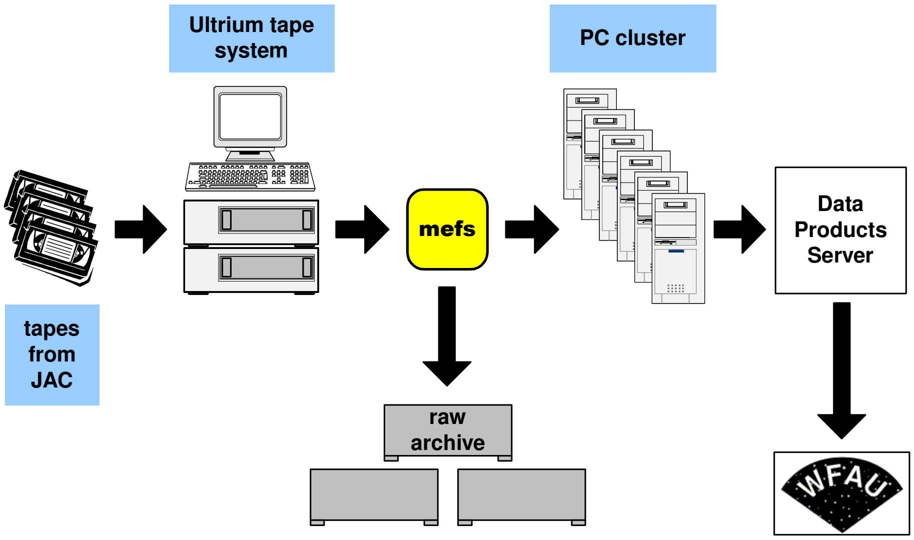

Figure 3. A schematic of the data processing setup for WFCAM in Cambridge. Incoming tapes are ingested, the data verified, converted to MEF format and fed to the processing cluster. A raw data archive is held on-line, whilst processed products are automatically transferred via the internet to the Wide Field Astronomy Unit (WFAU) in Edinburgh

Read combinations: data from all reads that comprise an exposure will be combined together to form a single frame.

Linearity correction: by and large the data obtained from NIR arrays are strongly non-linear. The linearity curve can be determined through observations of a stable "light source" for a range of exposure times. This can be done during the day, and it remains to be determined how often this will be necessary. Because each detector will be read out in channels (32 for each WFCAM detector and 16 for each VISTA detector) with each channel having its own electronics, it is very likely that a separate linearity correction will be needed for at least each channel. The linearity correction can be applied on-the-fly in the DAS or as the first reduction step in the main pipeline.

Each option has certain advantages and disadvantages. The former is more complex from an operational point-of-view since the non-linearity measures will need to be computed using a different readout mode to that in normal use and then fed back to the controlling system for application. Leaving the non-linearity operation for the pipeline simplifies the operational aspects but makes various (not unreasonable) assumptions about the timing of reset and readout operations and the stability of the illumination. One caveat is that pipeline correction of readout modes using multiple non-destructive reads would be very difficult, if not impossible.

3.2. Main Reduction Steps

In what follows we deal with the main reductions steps in their approximate location in the pipeline.

Linearity correction: if this has not been done in the DAS then it will be the first step in the main reduction pipeline (see section 3.1).

Dark and reset anomaly correction: dark frames (exposures taken with the light path blocked off) are used to calibrate out two separate additive effects:

- the accumlated counts that result from thermal noise (dark counts). This is generally a small, but not negligible effect.

- an effect, here called "reset anomaly", in which a significant residual structure is left in an image after the reset bias is removed in the DAS (see figure 4).

These effects can be removed together using a series of dark frames with an appropriate exposure time, and if needed, a simple scalable model for the reset anomaly to correct for illumination dependence.

Flatfield correction: because WFCAM and VISTA are multi-detector (and multi-channel) cameras, internal calibration will also involve accounting for the variation in mean gain from channel to channel. If the sky "on average" uniformly illuminates each detector then the variation in the "mean" counts on the flatfield map for each channel is a measure of the variation of the mean gain. Flatfielding is usually defined such that the "mean" counts in the object frame remain constant. For situations where there is a gain difference between detectors, the mean flat for each detector will in effect be normalised by the ensemble average counts over all detectors thereby ensuring correct inter-detector gain normalisation. The obvious caveat with defining differential gain corrections this way is that they are potentially a function of source colour due to inherent variation of the QE curves of the detectors (indeed this can even be the case at the individual pixel level).

Confidence map: the initial confidence map for each exposure is formed from the properties of the flatfields, and will therefore be the same for all images in a given flatfield sequence. Pixels outside a specified tolerance range are given a confidence of zero. Confidence maps are carried through further processing stages in conjunction with the "average" data frame background level noise variance. See section 3.3 for a fuller description of confidence maps and their use.

Defringe: if the fringe spatial pattern is stable and if flatfields can be generated without fringing present, it is possible to decouple sky correction and fringe correction and apply a defringing method similar to the one we have developed for optical imaging.[4] This involves creating a series of master fringe frames (eigenfringe maps) which are scaled by a suitable factor for each object frame. The scale factors are adjusted to minimise the fringe pattern in the processed frame.

Sky subtraction alternative: standard NIR processing often subtracts sky first and then flatfields. We can see why this can be advantageous compared with dark-correcting, flatfielding and sky-correcting by considering the following encapsulation of the problem,

$$D(x,y) = ff(x,y)\ [S(x,y) + F(x,y) + O(x,y) + T(x,y)] + dc(x,y) \tag{1}$$

where $D(x,y)$ is observed, $ff(x,y)$ the flatfield function, $S(x,y)$ is the sky illumination, $F(x,y)$ is the fringe contribution, $O(x,y)$ is the object contribution, $T(x,y)$ is the thermal contribution, $dc(x,y)$ is the dark current and without loss of generality we have excluded any explicit wavelength- and time-dependence for clarity. Stacking a series of dithered object frames with rejection produces an estimate of the terms

$$\hat{I}(x,y) = ff(x,y)\ [S(x,y) + F(x,y) + T(x,y)] + dc(x,y) \tag{2}$$

therefore,

$$D(x,y) - \hat{I}(x,y) = ff(x,y)\ O(x,y) \tag{3}$$

obviating the need for dark-correcting and fringe removal as both separate data gathering requirements and as separate data processing steps; and minimising the effect of systematic and random errors in the flatfield function by removing the largest potential error terms.

The caveats here of course are that this method may well remove parts of large extended objects, large area nebulosity, large low surface brightness objects and so on, unless suitable offset skies are used in the sky frame construction. Unfortunately this then opens the door for spatial and temporal variability of the sky background, leaving residual patterns.*

Reset anomaly: in the event that the dark correction stage above fails to remove the reset anomaly completely, then there could remain a challenging background variation over the detector to deal with. This is analogous to the problem of dealing with rapid variations in sky background level/structure during stacking/mosaicing. The reset anomaly in most current generation detectors appears to be a simple function of individual frame exposure time and general illumination level and appears to be amenable to modelling and removal using straightforward techniques.

Image persistence: images/artefacts from preceding frames can persist and be present on the current image. Strategies for dealing with this involve assessing the time decay characteristics and adjacency effects (*i.e.* image spreading). Correcting for image persistence will either involve updating and maintaining a persistence mask (for combination with the confidence map), or accumulating with suitable temporal decay, a persistence map, running over a night if necessary, to subtract from the current image.

Crosstalk: images from one detector channel may produce secondary images on other channels either on the same or on any other detector. In a stable environment, it is feasible to measure the contribution of crosstalk from one channel to another and thereby define a crosstalk matrix. This can then be applied as an additive correction.

Interleave: the on-sky detector pixel scale for WFCAM will undersample typical seeing conditions on Mauna Kea and in the very best seeing conditions the same will be true for VISTA on Paranal. To recover some of this lost resolution, an observation at a particular pointing may optionally consist of a number of microstepped exposures.[5] Microstepping is done by shifting the telescope a precise non-integral pixel distance, *e.g.* a 2x2 microstep sequence requires shifts of $n + 1/2$ pixels. Interleaving consists of creating an output image that is a regular interwoven pattern of **all** the input pixels, thereby sampling on a finer grid in an attempt to recover some of the lost resolution[6] The caveats are that interleaving preserves bad pixels (and indeed will generally affect more objects) and that the PSF may vary on short enough timescales to lead to unacceptably "spiky" interleaved PSFs requiring specialised analysis routines to deal with. Interleaving results in a new confidence map for the output image.

Dithering/Jittering: NIR detectors suffer from large numbers of bad pixels, cosmic ray hits and other cosmetic effects. In order to remove these (and to avoid saturation by the sky background), it is common practice to split a long exposure into several shorter exposures, which, rather than being repeated with each pixel looking at exactly the same sky position, are carried out at a series of different positions. This is similar to microstepping, but with image stacking usually at the same pixel scale. Combining dithered images with rejection allows for robust estimates of the sky and fringe patterns which are useful early on in the reduction procedure. After removal of the instrumental signatures, dithered images can be shifted and combined. The confidence map for each input image flags the bad pixels and a rejection algorithm is used to remove cosmic rays and other transient effects.

Astrometric Calibration: astrometric calibration is a multi-stage process and aims to provide each image with a World Coordinate System (WCS) to convert between pixel and celestial coordinates. This happens in the pipeline as follows:

Rough WCS: basic information in the FITS header, as well as knowledge of the instrument, allows us to give each processed frame a WCS that is accurate to within a few arcseconds. The basic

*For the sparse field situation with mainly point sources and small extended objects, any sentient object extraction algorithm with background following will do an excellent job at subtracting out smoothly varying background flux for each object at the extraction stage. Note that the background flux will be a combination of the sky emission, scattered light and diffuse astronomical emission (*i.e.* halo of a bright star or the disc of a galaxy). So whether the sky component is explicitly subtracted out or not, background following will still be necessary for accurate object extraction.

information needed is the RA and Dec of the pointing, a (stable) reference point on the detector grid for those coordinates (*e.g.* the optical axis of the instrument), the central pixel scale, the rotation of the camera, the relative orientation of each detector and the geometrical distortion of the telescope + camera optics, which defines the astrometric projection to use. Although most of this can be gleaned off-line, deriving accurate values through on-sky observations is an important task at commissioning time.

First Pass WCS fit: given a rough WCS for the processed frames, a more accurate WCS can be defined using secondary faint astrometric standards (*e.g.* USNO-B or 2MASS). The procedure is standard and straightforward: extract a catalogue for each detector; use the rough WCS to help match the (x, y) positions to the relevant subset of the astrometric catalogue; do an iterative clipped least-squares solution; update the WCS. For mosaicing (tiling) several of these frames together, this WCS now defines the transformation for mapping the input frames to the output grid.

Photometric zeropoint: For the purposes of quality control a photometric zeropoint will be determined for each observation by direct comparison of instrumental magnitudes with 2MASS. A more accurate photometric calibration is applied retrospectively given a complete night of observations including exposures in photometric standard fields (see 4.4).

Tiling: as the focal planes of both cameras are populated with detectors spaced between 42 and 94% of a detector size, it is necessary to take a series of exposures at large offsets to cover completely a single contiguous area of sky. Pros and cons of using such tiles as the basic unit of information are highlighted below.

pros: tiling makes better use of "fuzzy" dither edges around each of the detectors stacks in a contiguous tile. (A significant fraction of the total contiguous area will be on more than one detector stack.) Tiling generates smooth mosaic diagnostics and internal calibration statistics from combining many input detector stacks and also reduces by a large factor the later complexities of overlap cross-calibration. Finally, seamless coverage of large contiguous areas is a required data product.

cons: tiling generally involves non-linear resampling and hence use of sub-pixel interpolation schemes.[†] Sub-pixel interpolation affects the noise covariance matrix, complicating later processing stages. Finally, sky and seeing variations may make it difficult to achieve a "smooth" background on the output image and could also lead to problematic PSF variations.

An important aspect of tiling involves the choice of output WCS. We have catered for what we feel are the two sensible alternatives. For general ease of use one possibility is to use the TAN projection, with a reference tangent point defined to be the centre of the output image. The other possiblity is to keep the same projection used to describe the telescope+camera, for WFCAM and VISTA this is an ARC projection with Zenithal polynomial distortion. The disadvantage of the latter is the much higher radial distortion. Both lead to a problem with local flux conservation (see 4.4).

3.3. Confidence Maps

We define a confidence map as a normalised (to a median level of 100%) inverse variance weight map denoting the "confidence" associated with the flux value in each pixel. This has the advantage that the same map can also be used to encode for hot pixels, bad pixels or dead pixels, by assigning zero confidence. Furthermore, after image stacking the confidence map also encodes the effective relative exposure time for each pixel, thereby preserving all the relevant intra-pixel information for further optimal weighting.

The initial confidence map for each frame is derived from regular analysis of the master calibration flatfield and dark frames and is unique for each filter/detector combination due to the normalisation. To

[†]Note that in general stacking dither sequences may also involve interpolation if accurate pixel registration is not used, or if the optical distortion of the field-of-view is significant.

use the confidence maps for weighted coaddition of frames then simply requires an overall estimate of the average noise properties of the frame. This can readily be derived from the measured sky noise, in the Poisson noise-limited case, or from a combination of this and the known system characteristics (*e.g.* gain, readout noise).

All processed frames (stacked individual detectors, tiled mosaicked regions) have an associated derived confidence map which is propagated through the processing chain.

4. CATALOGUES

The standard catalogue generation software[7] makes direct use of the confidence maps for object detection and parameterisation producing quality control information, standard object descriptors and detected object overlay files. The possibly varying sky background is estimated automatically, prior to object detection, using a combination of robust iteratively clipped estimators. The image catalogues are then further processed to yield morphological classification for detected objects and used to generate astrometric and photometric calibration information.

Standard object descriptors include assorted aperture flux measures, intensity-weighted centroid estimates, and shape information, such as intensity-weighted 2nd moments to encode the equivalent elliptical Gaussian light distribution. The further processing pipeline adds PSF estimation and PSF fits for each object plus a generalised Sersic profile fit.

In addition to the object catalogue, the generation software produces a detected object ellipse overlay file to facilitate troubleshooting investigations via an image browser (*e.g.* DS9).

4.1. Quality Control

Every processed WFCAM and VISTA data frame will have a series of quality control measures automatically determined. These are designed to monitor the integrity of the instrument and also to provide measures that greatly enhance the usefulness of the data products for end users. These procedures occur in addition to routine daytime operational checks that monitor the general health of the system.

Quality control measures include night–time observing conditions and the status of the telescope + instrument detectors/filters/controllers and so on. This information is included in the FITS headers at the summits as part of the observing process.

Summit pipelines, and later the standard processing pipelines, will derive further quality control measures directly from the data and record them for later use. These measures will include: sky brightness – saturated? too bright? no signal?; sky noise – *rms* acceptable? pattern noise?; average stellar ellipticity – trailed images?; average FWHM seeing measure – for each detector, in focus?; stellar aperture corrections – weird PSF?; number of spurious images – corrupted data?; astrometric errors – wrong RA, Dec?; first pass photometric calibration – system throughput?; and on a nightly basis, derived local and global photometricity flags, limiting magnitude estimates, nightly zeropoints and extinction measures.

4.2. Morphological Classification

The object detection software produces a series of background-corrected flux measures for each object in a set of "soft-edged" apertures of radius $r/2$, $r/\sqrt{2}$, r, $\sqrt{2}r$, $2r$ ...... $32r$, where r is typically fixed as the median seeing for the site+telescope+camera. The average curve-of-growth for stellar images is used to define automatically an aperture correction for each aperture used and also forms the basis for object morphological classification.

The curve-of-growth of the flux for each object is compared with that derived from the (self-defining) locus of stellar objects, and combined with information on the ellipticity of each object, to generate the classification statistic. This statistic is designed to preserve information on the "sharpness" of the object profile and is renormalised, as a function of magnitude, to produce the equivalent of an $N(0, 1)$ measure, *i.e.* a normalised Gaussian of zero-mean and unit variance. Objects lying within 2–3σ are generally flagged as stellar images, those below 3σ (*i.e.* sharper) as noise-like, and those above 2–3σ (*i.e.* more diffuse) as non-stellar.

A by-product of the curve-of-growth analysis is the estimate of the average PSF aperture correction for each detector.

4.3. Astrometry Calibration

From the optical design studies of both WFCAM and VISTA we know that, to a good approximation, the astrometric distortion is well described by a radially symmetric polynomial distortion model.‡ The effective radial scale r' is related to the true scale r by $r' = r + k_3 r^3 + k_5 r^5$........ where the 5th order term is hardly discernible.

We intend, at least initially, to characterise the WCS using the ZPN projection[8] (*i.e.* ARC + polynomical distortion) using a 3rd order parameterisation. The coefficients for this are encoded in the FITS header using keywords like *e.g.* PV2_1 and PV2_3.

The effective scale due to the radial distortion is given by $dr'/dr = 1 + 3kr^2$, which for WFCAM and VISTA, respectively implies a scale change of ≈1.2% and ≈2.4% at the edge of the field-of-view compared to the centre. As an example, a 10 arcsec dither offset at the centre is therefore equivalent to a 10.12 and 10.24 arcsec offset at the edge (cf. central pixel scales of 0.40 and 0.34 arcsec respectively). For the outer layer of detectors this can lead to a large fraction of a pixel distortion across each detector during a dither sequence. In turn, this implies that non-linear resampling during stacking may be necessary.

In anticipation of this problem, we are implementing a range of interpolation schemes that offer a trade off between maintaining independent pixel noise and resolution degradation.

4.4. Photometric Calibration

The frame-by-frame catalogues provide the basis for the photometric calibration. An example use of this is for monitoring system throughput in almost real time to ≈5–10% accuracy using the 2MASS sources which will be detected in (virtually) every exposure. This will provide a first-pass photometric calibration and will be used to monitor sky transparency and system performance on minute timescales.

Since one of the goals of the science requirements is a survey-level calibration of 1–2% for both WFCAM and VISTA, a more accurate nightly photometric calibration will be determined from photometric standard fields observed hourly in each passband and will include measurements of overall extinction and zeropoints. We note however, that the internal gain correction, applied at the flatfielding stage, should place all the detectors on a common zeropoint system (at least to first order), and that given a stable instrumental setup, the apparent variation of zeropoint then directly measures the change in extinction without the need to rely solely on extensive standard field coverage over a range in airmass.

The standard fields will be set up during commissioning and the first year of operations. They are centred on previously defined primary calibration stars,[9] [10] and include sufficient stars ($>$ 100 per detector with J $<$ 16 chosen using 2MASS) that we will be able to measure spatial systematics and colour variation across each detector and the field-of-view of the instrument. This will enable direct quantification of the residual effects of scattered light correction and other distortions and enable accurate measures of the colour terms for each detector.

One of the contributors to the spatial systematics is the effect of the pixel scale change as a function of position. From the 3rd order distortion model the radial pixel scale $dr'/dr = 1 + 3kr^2$, hence the effective area on sky subtended by a pixel, which varies as $r'dr'$, changes as $\approx (1 + kr).(1 + 3kr^2)$. Flatfielding and mosaicing requirements are simpler to deal with by forcing/assuming "uniform" illumination of the background. A natural byproduct of this assumption is this additional term contributing to the photometric spatial systematics. However, since other more unpredictable factors such as scattered light will also play a significant role, it is simpler procedurally to bundle all the effects together and correct as one. A further complication caused by this is deciding how to trade off conserving flux in objects, which makes photometric calibration easier, and in sky, which makes mosaic tiling more challenging.

‡For more information see http://www.ast.cam.ac.uk/VDFS/

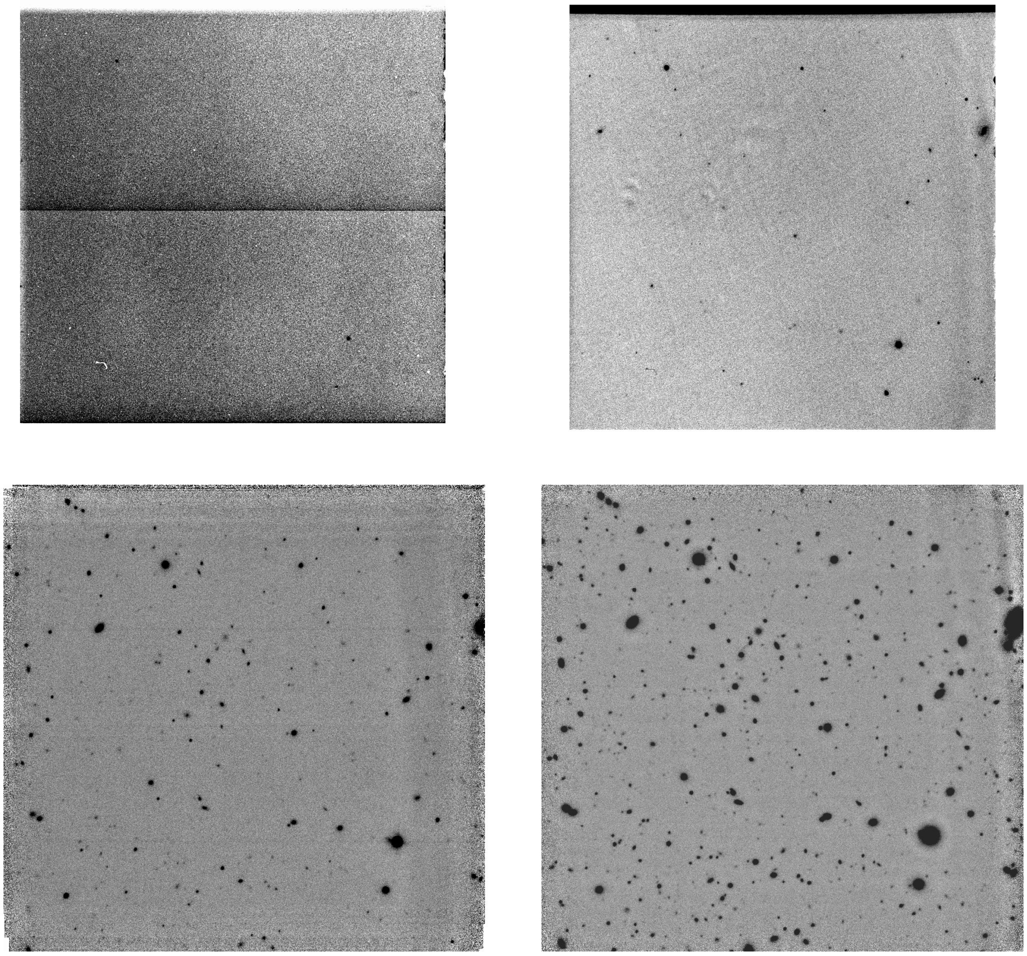

Figure 4. Top left: a raw 120s K_s-band frame; top right: the same frame after dark-correction, reset anomaly modelling, and flatfielding; bottom left: the result of pipeline processing and stacking a 40 frame dither sequence from a single observing block; bottom right the final K_s-band image stack of 1760 individual frames.

5. TRIAL COMMISSIONING

We have tested the WFCAM and VISTA pipeline modules on infrared data from several existing cameras. This includes the NIR ISAAC camera on the VLT. As a test of our pipeline and commissioning precedures, we downloaded the raw data for the 'Faint InfraRed Extragalactic Survey' (FIRES[§]) from the ESO archive. This survey consists of about 4000 science frames in J,H,K_s taken over a period of a year and centred on the Hubble Deep Field South, plus associated calibration data frames. These data have already been reduced and presented by Labbé *et al.* (2003)[11] and therefore make an excellent benchmark against which to test our reduction procedures.

[§]see http://www.strw.leidenuniv.nl/~fires

The reduction procedure follows very much the recipe outlined in section 3. The one major exception to this was that it was not possible to linearise the FIRES data as the necessary calibration data did not exist.

- Subtract an appropriate dark frame. This removed both the dark current and most of the reset anomaly.
- Divide by an appropriate flat field. These were formed by combining a series of images taken during sunrise and sunset.
- Create mean sky background frames by combining all of the images in a dither sequence with rejection. Subtract this estimate from each input image. We also found it necessary to create a temporally local residual sky estimate for each frame by averaging frames in a sliding window of size 5 centred on the image in question. These were then subtracted from the programme image.
- Combine each dither sequence into a single stacked image using cartesian offsets found from common objects on each frame. The combination makes use of the initial confidence map that was generated from the relevant filter flatfield image. A new output confidence map is generated.
- Generate a catalogue of objects on the stacked image.
- Using the catalogue and the initial WCS from the input header, calculate a WCS from the cartesian positions of objects in the catalogue and objects in the USNO-B astrometric catalogue.
- Classify each object in the catalogues as stellar, non-stellar or noise.
- Calculate a mean photometric zero point for each stacked image based on magnitudes in the 2MASS point source catalogue and instrumental magnitudes in the catalogues.
- Combine the stacked images in each passband from each observing block.

After this final step several low level artefacts were apparent, including an odd-even row modulation, low level charge anomalies, in rows, running through bright stars, and low level shadowing toward one edge and around the two brightest objects in the frames. In a similar manner to Labbé *et al.* , we reduced the impact of each of these effects using a robust, but ad-hoc, series of "filtering" operations on intermediate image stacks. Athough these operations greatly reduce the size of the effect they also remove any real large scale background variations. Although this is not a problem for this dataset this could be an issue in, say, Galactic Plane surveys.

A result of this testing is shown in figure 4. On the top left is a raw 120s K_s-band image. On the top right is the same image after dark subtraction, reset anomaly modelling and flat fielding. The obvious discontinuity between the top and bottom halves of the image have disappeared showing that the reset anomaly has been modelled out well and objects now appear over the sky noise indicating a good flat fielding result. The bottom right image is a stacked 40 point dither sequence from one observing block. Finally the bottom right panel shows the result of stacking 44 observing blocks (1760 individual images).

ACKNOWLEDGMENTS

The development of the WFCAM and VISTA pipelines is funded through grants from the UK Particle Physics and Astronomy Research Council.

REFERENCES

1. J. Emerson *et al.*, "Vista data flow system: overview," in *Optimizing Scientific return from astronomy through Information technologies*, *Proc. SPIE* **5493**, 2004.
2. N. C. Hambly *et al.*, "Vista data flow system survey access and curation: The wfcam science archive," in *Optimizing Scientific return from astronomy through Information technologies*, *Proc. SPIE* **5493**, 2004.
3. W. D. Pence, "New image compression capabilities in CFITSIO," *Proc. SPIE* **4847**, pp. 444–447, 2002.
4. M. J. Irwin and J. R. Lewis, "INT WFS pipeline processing," *New Astronomy Reviews* **45**, pp. 105–110, 2001.
5. M. M. Casali *et al.*, "The UKIRT IR wide-field camera (WFCAM)," in *The New Era of Wide Field Astronomy*, R. Clowes, A. Adamson, and G. Bromage, eds., *ASP Conference Series* **232**, pp. 357–365, 2001.
6. T. R. Lauer, "Combining undersampled dithered images," *PASP* **756**, pp. 227–237, 1998.
7. M. J. Irwin, "Automatic analysis of crowded fields," *MNRAS* **214**, pp. 575–604, 1985.
8. M. R. Calabretta and E. W. Greisen, "Representations of celestial coordinates in FITS," *A&A* **395**, pp. 1077–1122, 2002.
9. T. G. Hawarden *et al.*, "JHK standard stars for large telescopes: the UKIRT fundamental and extended lists," *MNRAS* **325**, pp. 563–574, 2001.
10. S. E. Persson *et al.*, "A new system of faint near-infrared standard stars," *AJ* **116**, pp. 2475–2488, 1998.
11. I. Labbé *et al.*, "Ultradeep near-infrared ISAAC observations of the hubble deep field south: Observations, reduction, multicolor catalog, and photometric redshifts," *AJ* **125**, pp. 1107–1123, 2003.