

Pipeline processing of wide-field near-infrared data from WFCAM (*draft - not for circulation*)

Mike Irwin^{1*}, Jim Lewis¹, Marco Riello¹, Simon Hodgkin¹,
 Eduardo Gonzales-Solares¹, Dafydd Wyn Evans¹, Peter Bunclark¹

¹ *Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, U.K.*

Received; Accepted

ABSTRACT

Over 30 Tbytes of raw WFCAM data from 3 semesters of observations have been transferred and pipeline processed by the astronomical survey unit in Cambridge (CASU). The resulting primary products are artefact-corrected images, interleaved and stacked as appropriate, and detected object catalogues. All pipeline products are routinely flux calibrated and astrometrically calibrated with respect to 2MASS. No further calibration is currently required since the pipeline output is within the overall requirements specification of 2% for JHK photometry and 100mas for astrometry. An assortment of quality control parameters (seeing, sky properties, limiting magnitude, overall stellar shape) are generated by the pipeline and are used to monitor and assess the quality and astronomical utility of the data and to provide input to a survey progress monitoring database. We show illustrative examples of the science quality of the data that is being generated and describe the technical solutions adopted to achieve these.

TBF.....

Key words: techniques: image processing - techniques: photometric - methods: data analysis - surveys - calibration:

1 INTRODUCTION

The UKIRT Wide Field Camera (WFCAM) on Mauna Kea with four HgCdTe Rockwell 2kx2k detectors on a 4m-class telescope, represents an enormous leap in deep NIR survey capability Casali *et al.* (2006). Each detector covers some $13.7' \times 13.7'$ of sky and the entire focal plane, with detectors separated by $\approx 95\%$ of the detector size, subtends a diameter of 0.9° . With nightly data-rates of a few hundred Gbytes (*i.e.* 10,000 individual $2k \times 2k$ images) automated pipeline processing and data management requirements are paramount. Furthermore, optimal processing of NIR data is far more technically challenging than for optical data, since NIR detectors are inherently more unstable and the sky (+thermal) background is over 100 times brighter than most objects of interest and varies in a complex spatial and temporal manner.

With a typical expected average data rate of an image of the sky every 5–30s, WFCAM 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 large amount 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 (*e.g.* Lawrence *et al.* 2006).

In order to deliver against various operational and scientific requirements, the WFCAM processing pipeline has to deliver astrometrically calibrated (to $\approx 100\text{mas}$) and photometrically calibrated (to $\approx 2\%$) science -quality reduced data. Furthermore, the data volume necessitates running most of the processing automatically, with in built-in quality control checks both of the data and of the pipeline procedures. The main exception to this is in the creation of some of the master calibration files, such as the flat-field frames, which are interactively created and updated at roughly monthly intervals.

Although several pipelines are involved in the end-to-end system: summit real-time processing; standard processing; and more advanced operations such as object profile fitting, here we focus on the standard processing carried out on WFCAM data in Cambridge. The summit processing is essentially a causal subset of the standard processing opera-

* E-mail: mike@ast.cam.ac.uk

tions, whereby the latter benefits from having a whole night of data available prior to commencement of processing and also has access to the latest calibration images (which may bracket the night being processed).

In this presentation we describe the pipeline architecture and algorithms developed to deal with the NIR imaging data from wide field cameras and illustrate the functionality by application to processing data from WFCAM¹, and discuss the main issues involved in creating 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 object catalogues.

2 DATA FLOW OVERVIEW

Although we are mainly concerned here with describing the pipeline architecture and software components developed to deal with the NIR imaging data from WFCAM, 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. By allowing selectable processing modules driven by recipe keywords in the FITS headers we can still retain flexibility within an automatic pipeline.

The general philosophy behind the pipeline processing is that all fundamental data products are FITS multi-extension files (MEFs) (Pence 2002) with headers describing the data taking protocols in sufficient detail to trigger the appropriate pipeline processing components. All derived information, quality control measures, photometric and astrometric calibration and processing details, are also incorporated within the FITS headers. Generated object catalogues are also stored as multi-extension FITS binary tables. These FITS files thereby provide the basis for incorporating information into databases both for archiving and for real time monitoring of survey progress and hence survey planning.

WFCAM is generally operated in either correlated double sampling (CDS) or Non-Destructive Read (NDR) mode and requires separate dark frames for each mode of use Casali *et al.* (2006). The WFCAM Data Acquisition System (DAS) produces reset-corrected frames and assembles each multi-sector read (4 quadrants of 8 channels per detector) into a coherent whole. The data is saved initially as four NDF files, one per detector, which are then converted to single channel FITS MEF format by the summit pipeline. The conversion is carried out independently for each detector on a dedicated machine. During the night a quality control summit pipeline, which uses a subset of the Cambridge processing operations, is used to provide (almost) real time assessment of critical observing conditions (*e.g.* seeing, sky brightness, image ellipticity) needed for the automatic scheduler and to help monitor telescope tracking and focus.

¹ The WFCAM processing pipeline is part of the VISTA Data Flow System development – see Emerson *et al.* (2004), Irwin *et al.* (2004) and Hambly *et al.* (2004) for an overview.

To reduce the data storage, I/O overheads and transport requirements, we make use of lossless Rice tile compression (*e.g.* Sabbey 1998) in all stages of the data flow system. For this type of 32-bit integer data, the Rice compression algorithm typically gives an overall factor of 3–4 reduction in image file size.

After conversion at the summit of Mauna Kea to Rice-compressed FITS files, data are shipped roughly weekly from the Joint Astronomy Centre using LTO tapes, one per detector channel. On arrival in Cambridge, on average about two weeks after the observations were taken, these tapes are verified² and the separate channels are combined to create the raw archived MEFs using Perl scripts and CFITSIO modules (Pence 2002). All raw and processed data is stored online using dedicated RAID5 or RAID6 disk arrays, each ranging from 4–10 Tbytes capacity. With Rice compression the yearly storage requirements for WFCAM are ≈ 20 Tbytes per year. The overall Cambridge pipeline setup is illustrated in figure 1.

After short commissioning periods in late Autumn 2004 and early Spring 2005, WFCAM has been in regular science use from April 2005. The main science program has been dominated by the UKIDSS suite of surveys (*e.g.* Dye *et al.* 2006), but also contains significant PI programmes. Since these programmes may be interleaved within a night, all WFCAM data is processed together in Cambridge.

Although somewhat variable, the average nightly data volume is over 2000 MEFs (≈ 150 Gbyte) of raw image frames, with peaks of up to 3600 MEFs, ≈ 230 Gbyte. Dark frames, with a range of different read-out modes and exposure times (to match the science frames) are taken every night and, on average, account for 100–150 image frames. Complete sets of twilight flat sequences are taken at roughly weekly intervals, several of which are combined at monthly intervals to form the master flatfield frames. The total amount of raw WFCAM data received and processed from the three UKIRT semesters (05A, 05B, 06A) to date, amounts to ≈ 30 Tbytes and corresponds to 168 raw LTO-I data tapes; all of which have been shipped and Aread without problem.

All the raw data is available online through the WFCAM raw data archive centre in Cambridge both for external users and also to facilitate reprocessing after pipeline improvements³. All the UKIDSS raw data and calibration frames are also transferred automatically to ESO, via the Internet, where they are stored in a dedicated archive accessible from the main ESO Science Archive Facility⁴. The data transfer rate to ESO averages ≈ 1 Mbyte/s.

Processing is carried out on a night-by-night basis using up to seven dual-processor 3GHz Xeon PCs, each with 2 Gbytes of physical memory and 850Gbyte RAID5 data buffers, running under a common Debian operating system. Longer term disk storage (all products, as well as raw images, are currently saved online) currently stands at

² There are occasional detector dropouts during observing, or detectors filled with NULL data values. Both of these are tested for during the verification and ingestion process and flagged for offline followup.

³ <http://archive.ast.cam.ac.uk/>

⁴ <http://archive.eso.org/>

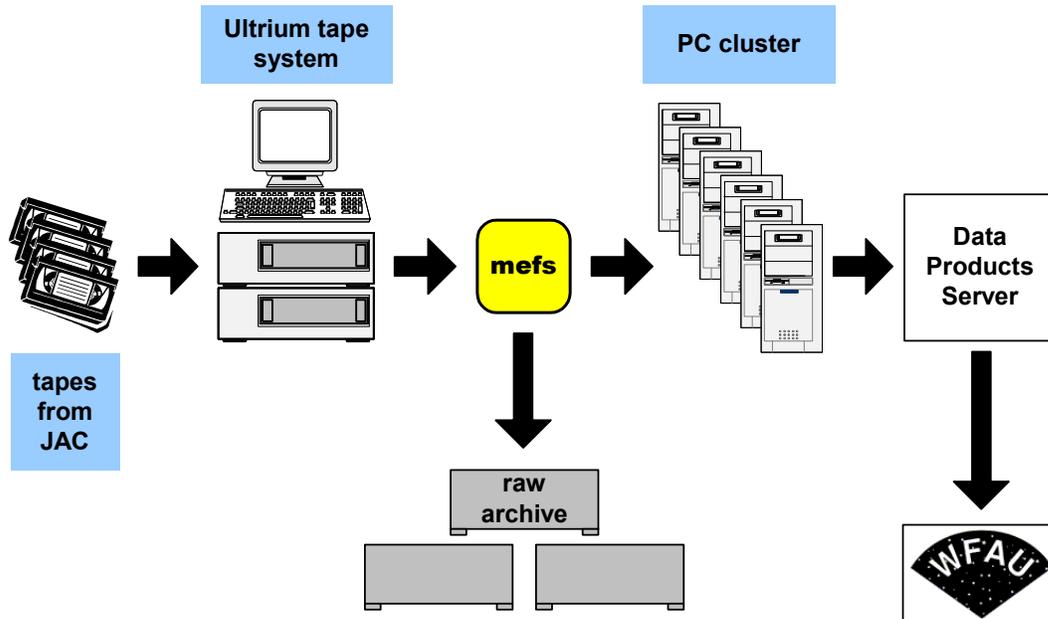


Figure 1. A schematic of the data processing setup for WFCAM in Cambridge. Incoming tapes are ingested, the data verified, converted to MEF format and fed to the processing cluster. A raw data archive is held on-line, whilst processed products are automatically transferred via the internet to the Wide Field Astronomy Unit (WFAU) in Edinburgh

~100 Tbytes of physical disk space and is via a mix of separate disk-server RAID5 systems and optical fiber-channel RAID5/6 arrays, all interconnected on separate G-bit backbones. As noted earlier, CASU store all image data, raw and processed, using lossless Rice compression which gives a factor 3-5 saving on disk space requirements.

Usually 8 nights are processed independently in parallel on 4 of the PCs. The processing time for a single night varies significantly, but generally depends on the average number of detected objects per field (*i.e.* crowded regions of the Galactic plane take longer) and the observing mode (interleaving and/or stacking). For an average night the total processing time is of order ~30 hours, but it can easily double, when for instance most of the data come from the UKIDSS Galactic Plane Survey (GPS) or Galactic Cluster Survey (GCS) that are observing in highly crowded regions.

The overall data processing strategy attempts to minimise the use of on-sky science data to form “calibration” images for removing the instrumental signature. By doing this we also minimise the creation of data-related artefacts introduced in the image processing phase. To help achieve this we make extensive use of twilight flats, rather than dark-sky flats (which potentially can be corrupted by thermal glow, fringing, large objects and so on) and by attempting to decouple, insofar as is possible, sky estimation/correction from the science images.

When the processing of a night is completed, various validation checks are made, *e.g.* were all processing steps carried out for each science file, are all calibration frames including sky frames present and so on? If these are satisfactory the image data products plus confidence maps are compressed for final storage and the quality information is ingested in a local Data Quality Control (DQC) database (see section 6). The ingestion process provides further checks on the presence of essential keywords and is a very powerful

way to spot obvious problems that may have occurred during processing. When the validation process is completed, the data is flagged as “ready to be transferred” to the WFCAM Science Archive (WSA) in Edinburgh (Hambly *et al.* 2006). The data transfer rate to Edinburgh can sustain ≈ 10 Mbyte/s which can readily keep up with the processed (compressed) data volume.

3 IMAGE PROCESSING

The processing pipeline can deal with compressed or uncompressed data in a transparent way because the I/O is managed through the CFITSIO library. However, since several I/O operations are required for each data frame, it is more efficient to uncompress images on-the-fly before processing commences, and then compress the final product for long-term storage.

Each night of data is pipeline processed independently using master calibration twilight flats (updated at least monthly) and a series of nightly generated dark frames covering the range of exposure times and readout modes used during that night. A running sky “average” in each passband is used for sky artefact correction. After removing the basic instrumental signature the pipeline then uses the header control keywords to produce interleaved and/or combined (stacked) image frames for further analysis. This includes generation of detected object catalogues, and astrometric and photometric calibration based on 2MASS (Skrutskie *et al.* 2006).

For orientation an overview of all the stages in the pipeline is shown in figure 2. In what follows we outline the general arithmetic operations required to remove the instrumental signature and then discuss the actual pipeline operations devised to deal with WFCAM data in more detail.

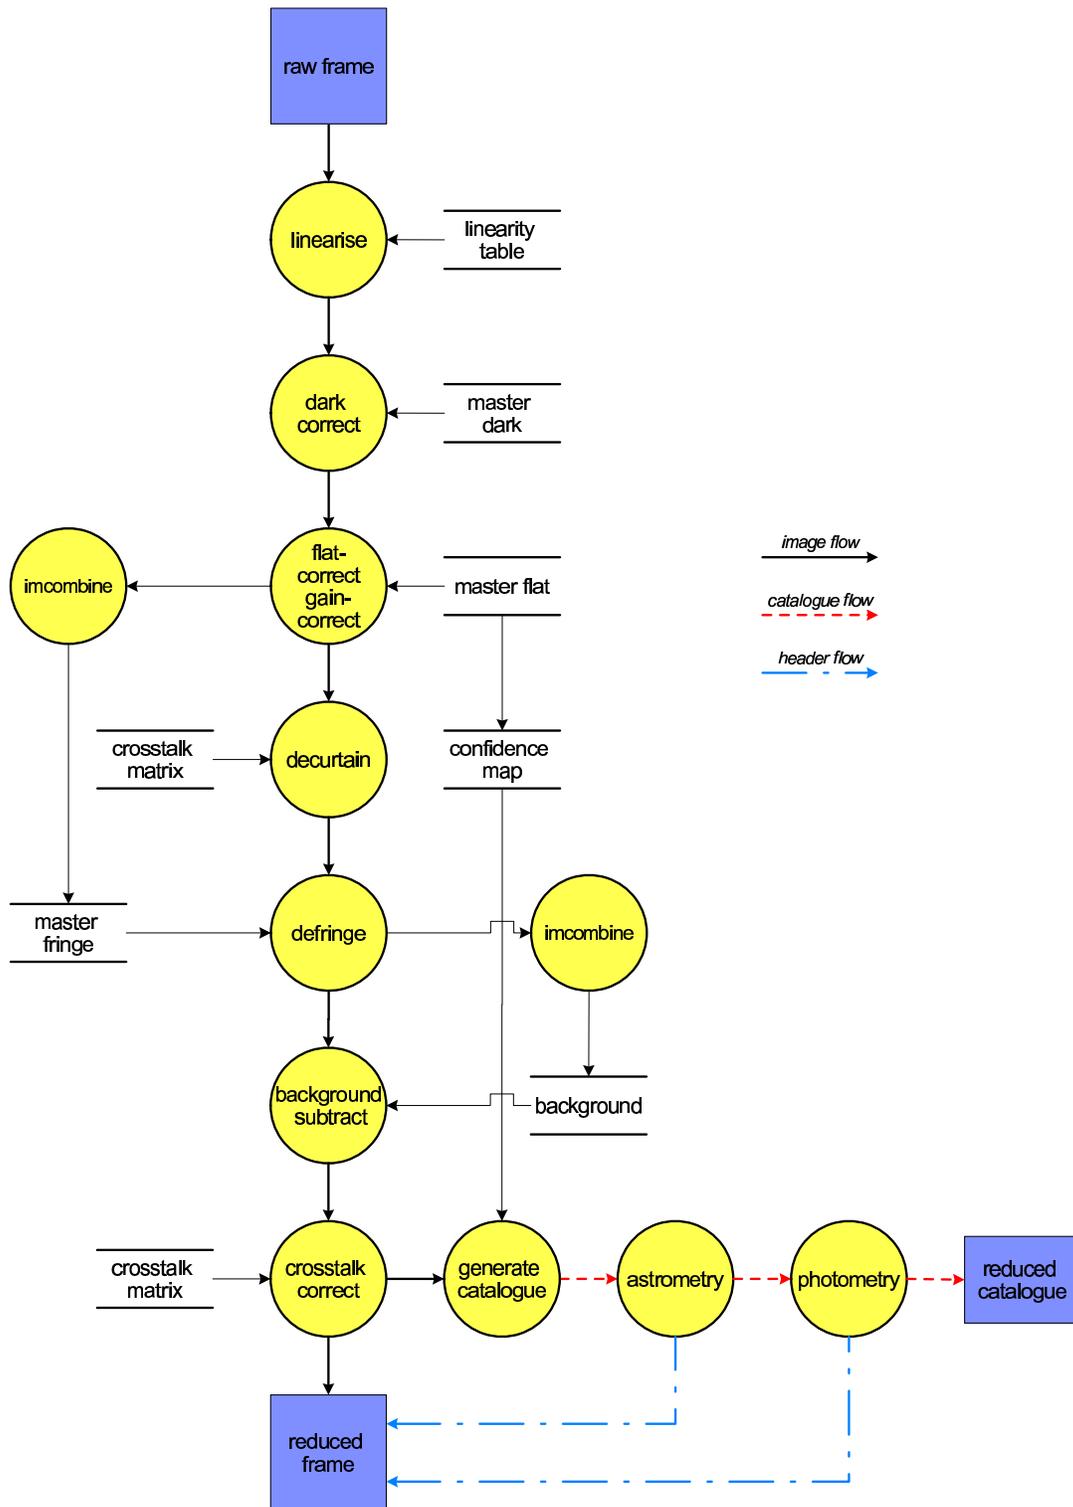


Figure 2. Overview of main stages in the processing pipeline for WFCAM.

3.1 Overview

Standard NIR processing recipes often subtract sky first and then flatfield. We can see why this can be advantageous compared with dark-correcting, flatfielding and sky-correcting by considering the following encapsulation of the problem

$$D(x, y) = f(x, y) [S(x, y) + F(x, y) + O(x, y) + T(x, y)] + d(x, y) \quad (1)$$

where $D(x, y)$ is observed, $f(x, y)$ is 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, $d(x, y)$ is the dark current (or more generally dark frame structure), and without loss of generality we have excluded any explicit wavelength- and time-dependence for clarity.

Stacking a series of dithered (jittered) object frames with rejection produces an estimate of the terms

$$\hat{I}(x, y) = f(x, y) [S(x, y) + F(x, y) + T(x, y)] + d(x, y) \quad (2)$$

therefore,

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

in principle 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. In the event that the dark correction stage fails to remove the reset anomaly completely, the residual background variation is analogous to the problem of dealing with short-term variations in sky structure and can be dealt with using the methodology above.

The caveats here, of course, are that this method may well remove parts of large extended objects, large area nebosity, large low surface brightness objects and more seriously, the photometry can be compromised by the presence of faint objects not rejected properly during the stacking and subsequent “sky” subtraction phase due to the data not being dark-corrected and flatfielded first.

The alternative is to treat the dark correction $d(x, y)$, flatfield $f(x, y)$, and fringe pattern $F(x, y)$, if present, as accurately known master calibration frames, in which case data processing involves solving the following variant of the problem

$$D(x, y) = f(x, y) [S(x, y) + k.F(x, y) + O(x, y) + T(x, y)] + d(x, y) \quad (4)$$

where k is a scale factor to be determined by the fringe removing algorithm, if required. In this case applying the master frames leads to

$$D'(x, y) = S(x, y) + O(x, y) + T(x, y) \quad (5)$$

reducing the problem to one of sky-correction, including residual instrumental effects such as reset anomaly, and then of detecting astronomical objects on an additive slowly spatially varying background. In this case stacking sequences of processed frames $d'(x, y)$ with rejection works much better, because the majority of the “noise” in the frames has already been removed prior to the rejection operation. Rejection of faint objects can be even further improved by iterating on this process.

The following sections provide, in order, more details for each of the processing operations required for WFCAM.

3.2 Pre-processing in the Data Acquisition System

To minimise the overall data volume several basic pre-processing steps are carried out in the WFCAM data acquisition system, including: reset-correction in both CDS and NDR mode; co-averaging successive exposures from within the same integration; and combining separate NDR mode reads into a single overall equivalent exposure.

The reset image (equivalent to a bias frame which is removed in situ) eliminates the need to write a separate image for each data frame and in principle should remove the need for further dark-frame corrections, as the dark current in these devices is essentially negligible. In practice this is not the case and dark frames are needed to minimise the residual “reset anomaly”.

3.3 Linearity correction

As a general rule the data obtained from NIR arrays may be strongly non-linear, although the linearity curve can be derived through observations of a stable “light source” for a range of exposure times (*e.g.* a sequence of dome flats). Potentially, because each WFCAM detector is read out in 8 parallel channels, for each of the four quadrants, 32 separate linearity correction functions may be needed for each detector.

In principle any 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.

We outline in Appendix A a general method for computing and applying corrections for non-linearity directly from CDS data. Although in principle for WFCAM this involves solving 16 million independent non-linear equations, one for every pixel, this is still technically feasible and we have devised a relatively simple and efficient scheme to implement this using Gauss-Seidel iteration.

Fortunately repeated dome flat sequences demonstrate that there is no significant WFCAM non-linearity (<1%) until close to saturation at 40K counts/pixel (though other artefacts such as channel boundaries appear at ~30K counts/pixel). An example of one set of measurements is shown in figure 3.

3.4 Dark correction

Darks are routinely computed from the daily observations by combining as many darks as are generally available (within reason) for each exposure time, number of co-adds and readout mode. For each detector the combining operation uses an additive correction to bring all darks in a stack to the same overall median background level, and then uses a 3-sigma

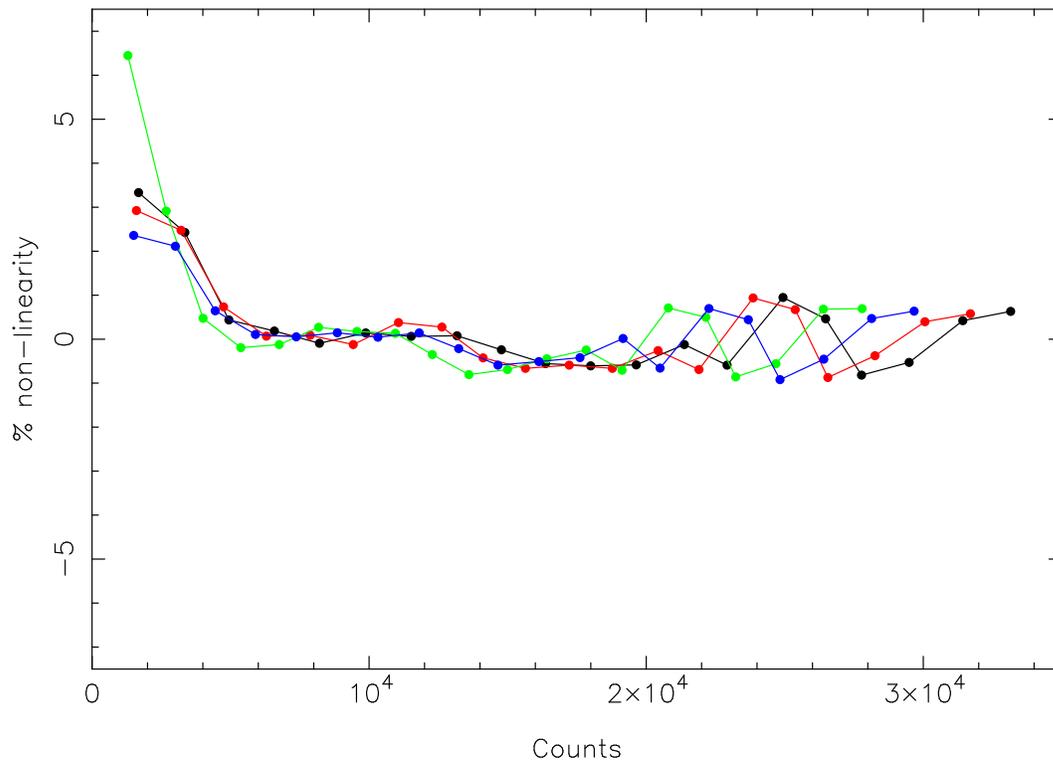


Figure 3. Linearity measurement from a WFCAM dome flat sequence. Each detector is separately colour coded. The apparent departures from linearity at high counts for each detector are highly correlated suggesting that stability of the illumination is the limiting factor here. At low counts a reset-anomaly pedestal DC offset dominates.

upper clipped modified median estimate for each pixel. The sigma clipping in this case is based on the average noise in the frame computed as $1.48 \cdot \text{MAD}$ where MAD is the median of the absolute deviation from the median (Hoaglin *et al.* 1983). The modified median is here defined as an average of the central two frames (if an even number survive clipping) and a 1-2-1 weighted average of the central three frames (if an odd number ≥ 3 survive clipping). This type of weighting ensures that the exact flux values in the original frames do not appear in the output frame, a well-known problem with straightforward medians.

If the required dark frame observations are not available, the nearest suitable calibration dark frame from a nearby night is used instead. If this still does not produce all the required darks to process a night of data, a suitably cunning combination of closely related dark frames is created and used instead.

In practice the dark frames are used to remove two separate additive effects: the accumulated counts that result from thermal effects - this is generally a negligible effect but occasional small regions suffer from significant dark current glow; and the reset anomaly which is manifest as a significant residual structure left on an image after the reset frame is removed in the DAS.

For WFCAM the maximum level of the reset anomaly reaches ~ 50 ADUs at the outer edges of the quadrants and the pattern is illustrated in figure 4.

After decurtaining (see sec. 3.8), the difference between successive dark frames can be used to estimate the global readout noise (RON) from differences of assorted dark frames. A comparison of single exposure CDS darks and

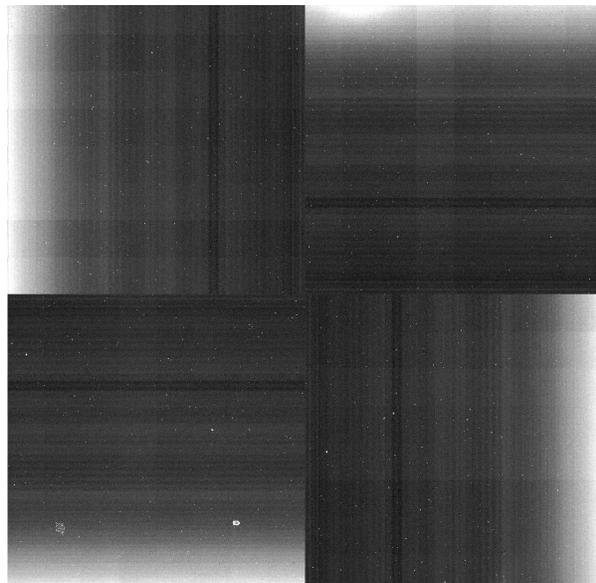


Figure 4. Combined 10s CDS dark frame for detector#1 for WFCAM. Note the exponential increase in the reset anomaly toward the readout amplifiers in each quadrant and the obvious hot pixels.

multiple read NDR darks is given in Table 1. There is no significant difference between the 5s and 20s CDS darks and only a small improvement in going to a 10-read 10s NDR

Table 1. Comparison of single exposure CDS darks and multiple read NDR darks.

Difference image noise				Exptime [sec]	Read Mode	<RON> [ADU]
#1	#2	#3	#4			
6.5	6.8	5.4	7.2	20	CDS	4.6
6.5	6.2	6.6	6.0	5	CDS	4.5
4.0	5.0	4.5	5.5	10	NDR	3.4

sequence.⁵ In NDR mode for WFCAM the output frame is computed internally in the DAS from the gradient of the flux increase for each pixel (*i.e.* flux/s) and then normalised back to the equivalent total counts for that exposure time. Long term monitoring of dark frame patterns for WFCAM indicate that darks are stable on time-scales of a few days or more.

3.5 Flatfielding

Weekly dawn twilight flatfield sequences are taken using series of 9-point jitter sequences in Y, H one night and Z, J, K the next. These are dark-corrected and then stacked to form intermediate master flats. The stacking follows the same procedure used in forming dark calibration frames with the exception that multiplicative, rather than additive, scaling of the median background levels is used to place all frames in a sequence on the same overall level prior to combining. A further refinement is to supplement the 3-sigma median clipping to enhance removal of faint objects by using a second pass through the data. That is, the first estimate of the flatfield is used to flatfield each input frame in the stack before repeating the 3-sigma rejection process. Since this operation typically reduces the effective sigma by a factor of ≈ 30 it is very successful in removing faint astronomical objects from the combination procedure.

If available, intermediate flats from two or three such sequences are then further combined, with the aim of achieving ≈ 1 million photons per pixel (*i.e.* $\approx 0.1\%$ photon noise) from the total combination. We find that these flats are stable on at least monthly timescales, give good dark sky correction (*i.e.* gradients are at the $\sim 1\%$ level at most), and show no fringing or measurable thermal emission. The overall QE is good, but the flats show large spatial gradients across the detectors with up to a factor of $\times 2$ sensitivity variation. Since the general characteristics of these variations can be seen across all filters, the simplest interpretation is that these variations in level reflect genuine sensitivity variations across the detectors (this has been subsequently confirmed by examining χ^2 values of residuals from PSF-fitting).

As an illustration, all four detector H-band master flats for the period 7th–19th April 2005 are shown in the first row of Fig. 5. The sensitivity variation on all detectors is reflected in the confidence maps (see section 3.13) and directly impacts the uniformity of achieved survey depths. This is quantified in the series of histograms, second row of

Table 2. RON and gain estimates based on April 2005 WFCAM data.

Detector	Dark DC level [ADU]	RON [ADU]	Gain [e^- /ADU]	RON [e^-]
#1	-2.4	3.8	4.84	18.4
#2	3.7	4.2	4.87	20.5
#3	-9.0	4.2	5.80	24.4
#4	49.1	4.4	5.17	22.7

Fig. 5, which show the recorded sky level normalised to the median level of all four detectors (*i.e.* these are notionally the gain corrections). Although the “average” sensitivity is good, there are significant regions on most detectors a factor of 2 worse, these appear as darker regions in the first row of Fig. 5. These variations are also present in the dark sky (*i.e.* they are not caused by weird illumination gradients in twilight flats) as evidenced by the uniform background seen in flatfielded dark sky data.

3.6 Gain correction

As WFCAM is a multi-detector (and multi-channel) camera, flatfielding also involves an internal system calibration which accounts for the variation in mean gain/sensitivity from detector-to-detector (and channel-to-channel if required). Implicit is the assumption that the sky “on average” uniformly illuminates the focal plane and thence that the variation in the “mean” counts of the flatfield measures the variation of the mean gain.

In the case of a single detector camera the “mean” of the flatfield image is usually normalised to unity. This ensures that when the flatfield correction is applied the average counts in the output image are the same as for the input image. For multi-detector camera such as WFCAM we normalise by the ensemble average counts over all detectors thereby ensuring correct inter-detector gain normalisation. The advantage is that only a single overall external calibration is required for each passband. The obvious caveat with defining differential gain corrections this way, is that they may be potentially a function of source colour due to inherent wavelength-dependent variation of the QE curves between the detectors (*e.g.* the twilight sky is a different colour from the dark sky, although tests of twilight flatfielded dark skies reveal no measureable effect from this).

3.7 Inter-pixel correlation and gain estimates

Although not strictly part of the pipeline processing we briefly discuss here for completeness our measurements of the gain (e^- /ADU) and inter-pixel correlation of the detectors.

Estimates of the gain (and RON) are regularly derived using sequences of darks and dome flats. Table 2 provides an illustrative example of these measurements and highlights the level of the global reset anomaly pedestal level seen in dark frames and also the derived “apparent” gain in e^- /ADU. The gain here is the average of values measured over each detector at three different background levels, 23k, 14k and 5k counts. The overall variation in measured gain is

⁵ Note that N NDR frame reads for $N \gg 1$ asymptotically give $\approx \sqrt{N/3}$ improvement (not \sqrt{N}) in RON over normal CDS mode because the resulting readouts are correlated.

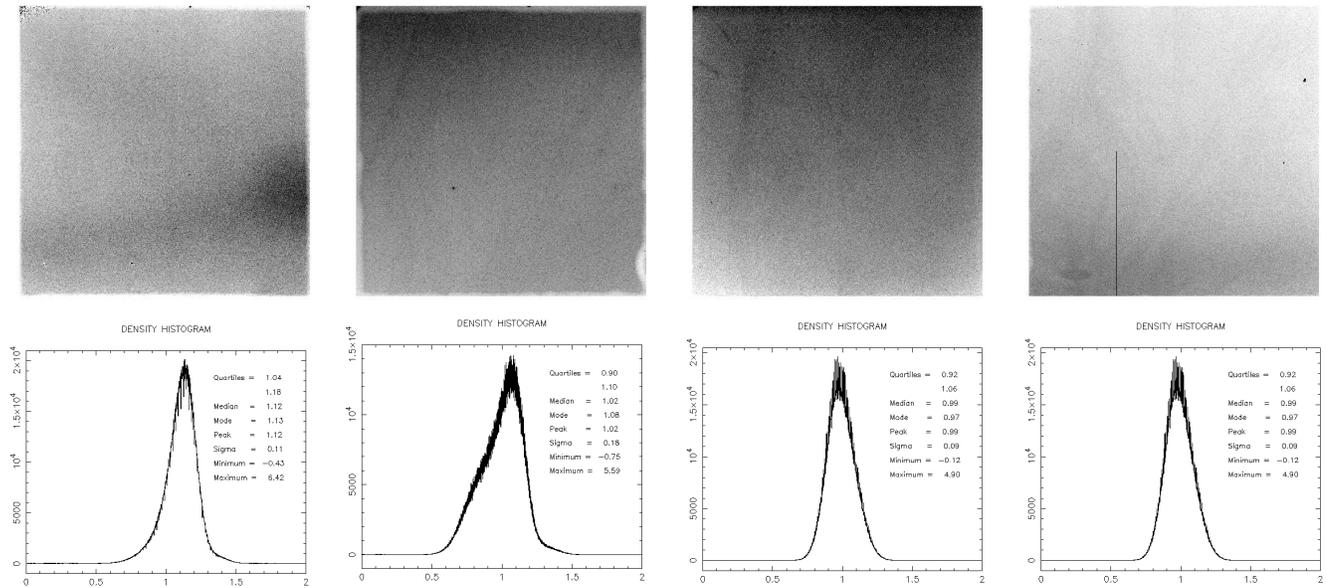


Figure 5. From left to right chip #1 to #4. First row: H-band master flatfield image; second-row: normalised histograms of the flatfield frame levels showing the overall variation in sensitivity.

at the $\sim \pm 0.05$ level with no clear trends with background level.

Whilst doing these measurements of the gain from the dome flat sequences, we could also use the same dataset to make an estimate of the inter-pixel capacitance via a robust measure of the noise covariance matrix. This requires a high photon (noise) level and the difference between two dome flats taken under constant illumination levels provides a simple way to generate this. All four detectors give remarkably consistent results and the central part of the normalised noise covariance matrix, $c_{i,j}$, defined by

$$c_{i,j} = \frac{\langle d_{k,l} d_{k+i,l+j} \rangle}{\langle d_{k,l} d_{k,l} \rangle} \quad (6)$$

where $d_{k,l}$ is the difference image, is summarised below

$$C = \begin{bmatrix} 0.01 & 0.04 & 0.01 \\ 0.04 & 1.00 & 0.04 \\ 0.01 & 0.04 & 0.01 \end{bmatrix} \quad (7)$$

with essentially zero coefficients in adjacent rows and columns further out. Integrating over the noise covariance matrix, C , close to elements 0,0 gives a value of 1.20 which implies that the total reduction in directly measured noise variance (*i.e.* that measured using a conventional method) due to the implicit “smoothing” is therefore also 1.20 (from application of Parseval’s Theorem and the power spectrum \leftarrow FT \rightarrow autocorrelation duality). This therefore predicts a $\sim 20\%$ overestimate of the gain, and hence a $\sim 20\%$ overestimate of the QE coefficients. This behaviour has been more directly measured for Rockwell Hawaii-II detectors by (Finger *et al.* 2006).

3.8 Decurtaining

After dark-correcting and flatfielding, a pseudo-periodic ripple is still present on all frames at the ± 5 ADU level. This is removed by exploiting the 4-fold quadrant symmetry of the effect (see figure 6) to compute a bilinear correction array. In particular, the algorithm assumes that the variation is the same in each quadrant of the detector and then proceeds to compute a quadrant level row-by-row and column-by-column correction based on an iteratively k-sigma clipped median for each row and column. The confidence map (see section 3.13) is also used to flag and avoid using bad pixels for this.

The column and row arrays are then filtered using a 5-point running median followed by a 3-point running box-car filter; normalised to have zero median; and then subtracted from the input image. An illustration of the curtaining problem and its removal is shown in figure 6.

Occasionally a single-channel pedestal offset is seen, particularly on one quadrant of detector #4. This shows up at the several ADU level but is left un-corrected. We also note that the decurtaining algorithm also corrects for the majority of any 4-fold symmetric leftover reset anomaly. Furthermore, the non-astronomical symmetry, and the use of a robust iterative median estimator, help ensure that the decurtaining algorithm is insensitive to the presence of real astronomical sources and even works well in crowded Galactic plane regions.

3.9 Defringing

Atmospheric emission lines may cause interference fringes to be present in the sky background, typically at the level of a few % of sky. Since the fringes can have complex spatial structures on a range of physical scales on the detector, removing them successfully is usually a multi-stage process. Although WFCAM does not suffer from significant fringing,

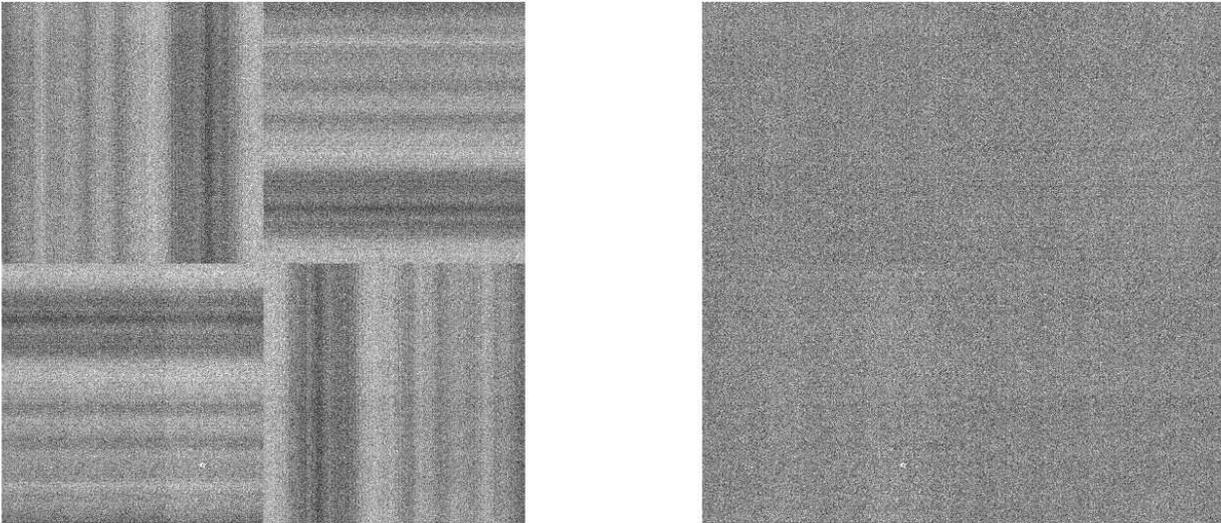


Figure 6. Curtaining problem, illustrated by taking the difference of two dark frames of the same type (left panel). The right panel shows the same image after applying the decurtaining algorithm.

we outline here, for completeness, the procedure we have developed to deal with it.

First we note that fringing is an additive effect, so if removed as part of dark sky self-calibration, this would introduce a systematic error in the photometry. To effectively perform sky fringe removal requires the flat fielding to be decoupled from the defringing by, for example, using twilight sky exposures to construct the flatfield frames, where the contribution from sky emission lines is negligible.

Consequently, the first stage of the process is to correctly flatfield the dark sky science data and use a sequence of offset sky exposures to construct a fringe frame. These input frames are combined after suitable scaling to match the background levels and sigma-clipping to remove astronomical objects.

The defringing process then requires solving for the fringe scale factor k in equation 4. Since the fringe pattern is characterised by more rapidly varying spatial structure than the sky and thermal contributions, the overall background variation on the target and fringe frame is temporarily removed by use of a robust low-pass filter such that

$$D'(x, y) \approx k.F(x, y) + O(x, y) \quad (8)$$

The objects are localised, therefore a simple overall robust background noise estimator based on the median of the absolute deviation from the median ((Hoaglin *et al.* 1983), can be used iteratively to find the scale factor k that minimises the background noise in $D'(x, y)$. Allowing the scale factor to vary ensures that the the relative contribution of the sky emission lines, which may vary in strength, is correctly dealt with.

Of course, if suitable sky-correction frames can be constructed then, in general, defringing will be automatically taken care of in the sky subtraction phase, discussed in the next section.

3.10 Sky subtraction

We have experimented with several alternative sky subtraction strategies ranging from computing a single sky correction for each band per night, to using the data within a single observing block to construct a local, *i.e.* spatial and temporal, sky frame. Scattered light within the camera, which varies as a function of position on the sky (see for example Dye *et al.* 2006), and illumination and exposure time dependent artifacts in the system, preclude using a single band-dependent sky correction for the night.

We found that grouping the sky estimation and correction stage by exposure time within each passband was necessary due to the presence of a combination of additive and multiplicative artefacts in the sky correction frames. The generally more spatially localised multiplicative artefacts scale closely with exposure time whereas the additive components do not. The latter includes residual illumination-dependent reset anomaly and pedestal offsets.

Naturally, computing reliable sky correction frames from within an observing block, *i.e.* from within a usually small patch of sky, is not always feasible, *e.g.* when observing close to the Galactic Plane. To overcome this the pipeline groups images temporally, and as well as spatially. If local sky estimation is not possible, sky frames are created by combining several adjacent sets of science exposures with sufficiently large offsets to ensure a proper rejection of all astronomical sources. The sky frame is then scaled to the sky level of the image to be corrected, its median level is adjusted to be zero and it is then subtracted off, thereby leaving the average sky level in the image the same.

Best results are obtained by producing different master sky frames for each filter/exposure-time/elapsed-time combination. To achieve this, the pipeline first groups the science frames by filter and then, for a given filter, by exposure time. After that, for each sub-group as many lists as are compatible with a fixed minimum and maximum number of frames

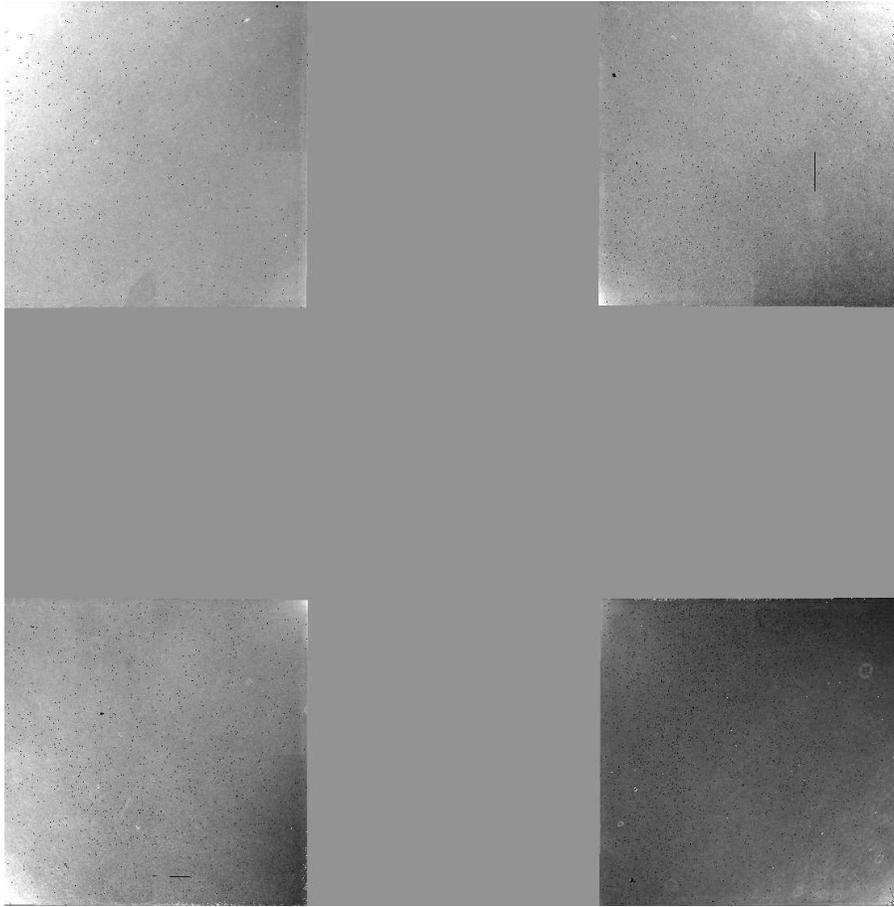


Figure 7. A K-band sky correction frame aligned using an inherited WCS and displayed in the conventional orientation, with North to the top and East to the left, showing the correct relative detector locations on sky. A number of artefacts are visible due to a combination of scattered light off dusty surfaces, particularly in the bottom right of the picture, and negative (darker) regions mainly due to bad pixels.

and with a minimum angular separation on the sky (to ensure proper object rejection) are formed.

Forming a master sky frame is a two-stage process. First, potential sky frames within a dither offset, or microstep sequence, are combined using a similar recursive methodology to that used for creating the master dark and flatfield frames. At this intermediate stage, because the dither and/or microstep offsets are only a few to several arcsec, some residual structure from astronomical objects may remain. As an illustration of the robustness of the combining procedure, figure 8 shows the result of combining 5 coaveraged 30s J-band exposures taken as part of a survey of M33, with dither offsets a combination ± 6.4 arcsec from the reference position. Although low level artefacts due to the presence of bright stars are still just occasionally visible the majority of the astronomical sources are effectively removed, even at this first stage. To ensure a more comprehensive removal of astronomical residual artefacts, groups of several such intermediate products, chosen as described earlier, are then further combined in a similar manner, to create the master sky-correction frames. We find that this double non-linear iteratively clipped median combination is extremely robust against the presence of even faint astronomical sources.

On average, a typical night will have between 2 and 10 master sky frames per filter, but in some extreme cases as many as 30 per filter can be produced. We show an example set of sky correction frames for each detector in figure 7.

3.11 Cross-talk

Cross-talk artefacts (*i.e.* pickup in adjacent channels) are confined within detector quadrants and occur between the 8 channels readout in parallel in each quadrant. All the detectors show similar cross-talk patterns with the induced artefacts being essentially spatial derivatives of saturated stars, with either a “doughnut” appearance from heavily saturated regions (see figure 10) or half-moon-like (*ie.* positive/negative residual) appearance from only weakly saturated stars. These occur at integer multiples of ± 128 pixels (*i.e.* multiples of the the channel width) either side of the saturated star. The cross-talk images are symmetric (quadrant boundaries notwithstanding) with respect to the source star and typically induce features at $\approx 1\%$ of the differential flux of the source, dropping to $\approx 0.2\%$ and $\approx 0.05\%$ further out (see figure 11). Beyond three adjacent channels the effect is generally negligible. Figure 9 illustrates where cross-talk images are expected in detector X-Y space, relative to

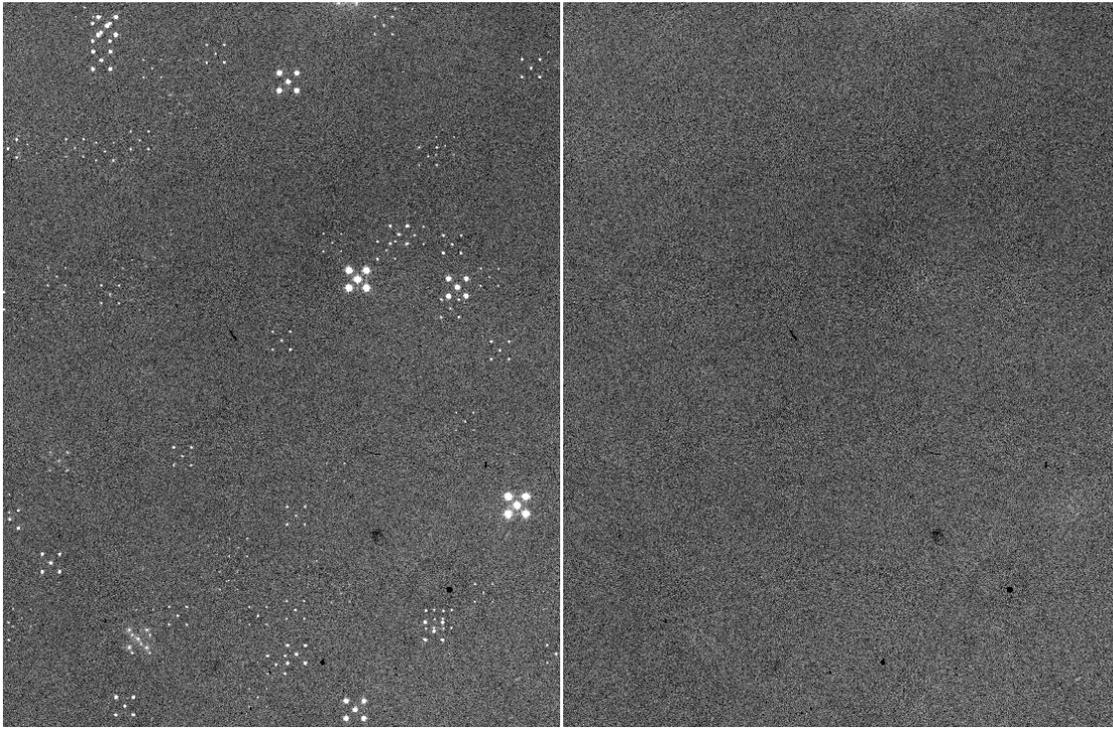


Figure 8. An illustration of the robustness of the combinatorial operations used to form sky frames; left panel: a $10' \times 8'$ region of the average of the 5 images to be combined, taken from a ± 6.4 arcsec dither sequence; right panel: the result of the combination.

the position of saturated objects for each quadrant (note that the orientation of the detectors on-sky also rotates by 90 degrees within the detector mosaic array). The rotating cross-talk pattern within a detector arises because the read-out amplifier for each quadrant is on a different edge of the detector (denoted by the red lines in the figure).⁶

The edge cross-talk seems to be the directional derivative of the signal in any one channel and is most strongly coupled to adjacent channels, but also appears at decreasing levels throughout the quadrant. It does not however cross quadrant boundaries and is only measurable for saturated or close to saturated objects. We don't understand the detailed cause but it is unlikely to be a direct property of the controllers.

To fix it we construct a model of the expected cross-talk for saturated or close to saturated images and compare the model with the regions $\pm nx128$ pixels either side of the

⁶ This type of cross-talk was unexpected. All other (optical) mosaic cameras we have encountered produce secondary images (ghosts) on adjacent channels and/or detectors which are miniature scaled +ve or -ve versions of the original image. For a stable system, it is feasible to measure the contribution of cross-talk from one channel to another using bright point-like sources and define a comprehensive cross-talk (generally asymmetric) matrix $C_{j,k}$. Providing the cross-talk terms are small (*i.e.* $< 1\%$, the most likely scenario), a simple single-pass additive correction scheme can be used to correct for the effect, $I'_j = I_j - \sum_{k \neq j} I_j C_{j,k}$ where I_j is the observed frame and I'_j the corrected version. The typical error in making a single pass correction is approximately $< C_{j,k} >_{j \neq k}^2$, which governs the requirement on the magnitude of the cross-talk terms.

source of the cross-talk. After some experimenting we ended up using the derivative of the flat-fielded data to form the model since the raw data gave too noisy a model (unsurprisingly). There's a couple of free parameters - the coupling coefficient; and a scaling parameter relative to saturation (*i.e.* unity for saturated parts and rapidly tapering to zero well-away from saturation). The comparison with the real data effectively produces a clipped combination of model and real cross-talk images (if possible) to form the "final" correction to be applied.

It doesn't remove the artefacts completely but does reduce them by about a factor of 10.

Creating a correction for the WFCAM derivative-like artefacts has involved devising a robust combination of a model involving the directional derivative of the primary source, and if possible, comparing and combining this with the actual cross-talk artefacts to create a correction sub-image. This extra artifice is required because the model depends on noisy derivatives which are also only an approximation to the real effect. By careful construction the model sub-image can be made immune to the presence of real objects overlapping the artefact and although the derivative is only an approximation, a relatively clean subtraction generally results.

The original artefacts are generally non-astronomical in appearance and so are any remnant residuals after cross-talk correction. Although cross-talk images do not straddle the detector quadrant boundaries, or occur between detectors; unfortunately, they do not generally stack out when combining dither sequences. (*e.g.* figure 11).

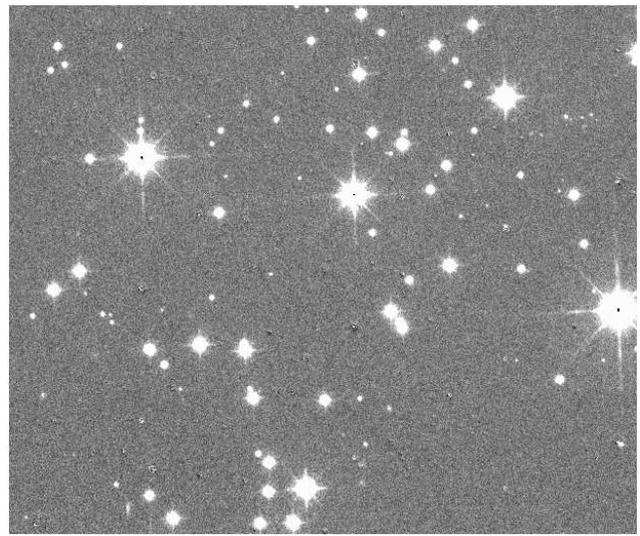
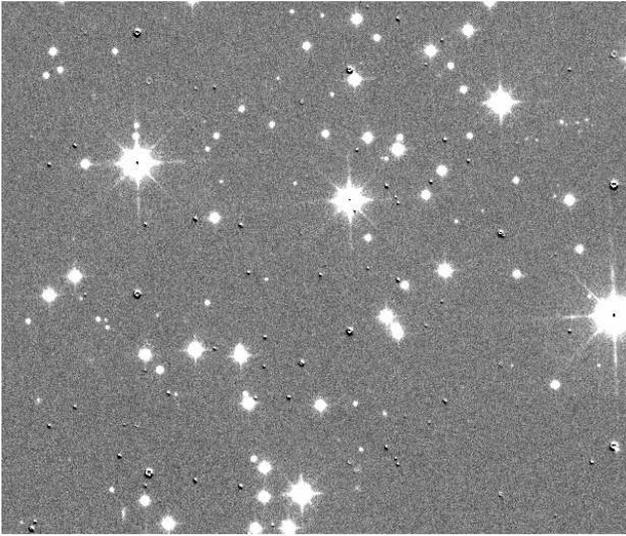


Figure 11. Left panel: an example of a stacked dither sequence image centred on the open cluster M67 highlighting several aspects of the cross-talk problem. Right panel: the residuals from cross-talk artefacts can be reduced by a factor of ≈ 10 but not completely eliminated.

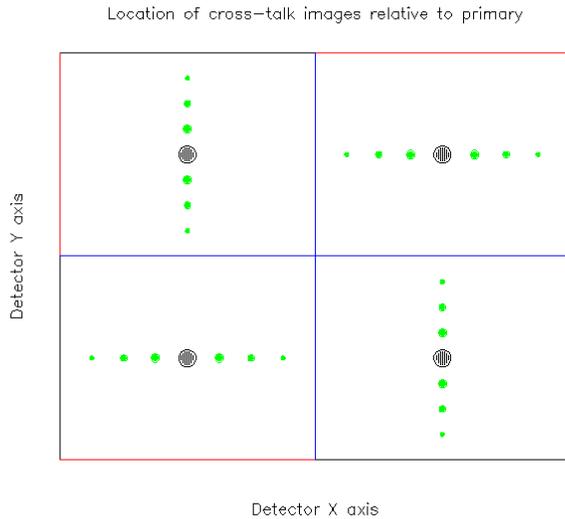


Figure 9. Representation of the cross-talk pattern as it appears on each of the four quadrants of each detector. The red lines denote the position of the readout amplifier for each quadrant.

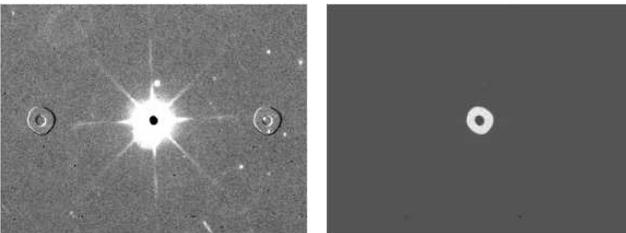


Figure 10. Cross-talk from a heavily saturated image (left panel), with an alternative view of the saturated image (right panel) to illustrate the cause of the problem (*i.e.* the cross-talk image is, to first order, the derivative along the readout direction of the signal in the saturated image channel).

3.12 Image persistence

Astronomical images, and artefacts from preceding frames, can persist and hence be present on the current image. Strategies for dealing with this involve assessing the time decay characteristics and adjacency effects (*i.e.* image spreading) if present. In the case of no image adjacency effects, correcting for image persistence should in principle 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. For example, in the simplest case with no image adjacency effects we might expect

$$I_k^{obs}(x, y, t) = I_k^{true} + f \times I_{k-1}^{obs}(x, y, t - \Delta t) \times e^{-\Delta t/\tau} \quad (9)$$

where k is the image sequence number, f is the instantaneous fraction of the image persisting after frame reset(s), Δt is the time interval between frames, and τ is the persistence decay constant. However, what we have found for WFCAM data is that the level of persistence is to a large part unpredictable. Although the temporal decay is fairly well understood and predictable with an e-folding time of ≈ 40 s, the fractional level of the persistence, f , which is typically $\approx 0.1\%$, shows an unexplained scatter almost as large as the effect itself, which precludes using a simple fix.

The alternative strategy of flagging regions potentially affected by persistence artefacts is a possible solution, but at this stage it is unclear whether this is necessary to deliver the main science goals. For programmes involving stacking, persistence images are rarely a problem since they are effectively removed, or reduced to negligible level, during the stacking/rejection process. Figure 12 illustrates the effect and shows an example of worse-case persistence from a heavily saturated star.

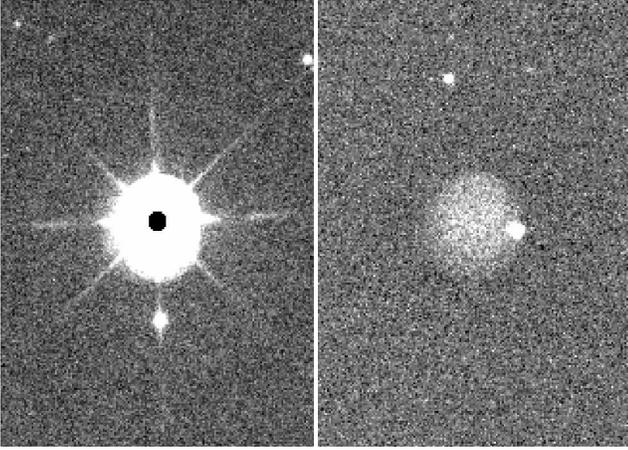


Figure 12. An example of worst-case persistence from a 6th magnitude (2MASS) star; left panel 10s H-band exposure; right panel 10s H-band exposure taken 15s later on an adjacent field. The level of the persistence plateau is ≈ 100 counts above sky (12000 counts) *i.e.* at a surface brightness of about 19.5 mag/sq arcsec.

3.13 Confidence maps

We define a confidence c_{ij} map as a normalised (to a median level of 100%, *i.e.* $\langle c_{ij} \rangle_j = 1$) inverse variance weight-map denoting the “confidence” associated with the flux value in each pixel j of frame i . This has the advantage that the same map can also be used to encode for hot, bad 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 inter-pixel information for further optimal weighting.

For WFCAM the initial confidence map for each frame is derived from regular analysis of the master calibration flatfield and dark frame sequences and is unique for each filter/detector combination due to the normalisation and global detector/filter properties. We define bad pixels in the confidence map in two stages. First we compare (at least) two independent sets of stacked dark-corrected twilight flatfield frames, with average count levels differing by a factor of roughly two. Pixels that do not flatfield from one set to the other within 4-sigma of the ensemble average are flagged as bad. In addition, we analyse the master flatfield frame to look for pixels with significantly different response (*i.e.* 25% better or worse) from their local neighbourhood average. To achieve this we essentially correct the master flatfield map first for any global illumination and/or sensitivity gradients using the same algorithm we use for background following in the object cataloguing stage (see section 4). Finally, the remaining good pixels are assigned a confidence according to their overall sensitivity, as deduced from the input master flatfield. These primary confidence maps are updated on the same timescale as the master flatfields.

To use the confidence maps for weighted co-addition of frames, or for object detection, then simply requires an overall estimate of the average noise properties of the frame. This can be readily 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, RON). All processed frames have an associated derived confidence

map which is propagated through the processing chain in the following manner.

Defining the signal s_i in frame i with respect to some reference signal level s_o as $s_i = f_i s_o$, where f_i denotes the relative throughput (which in photometric conditions would be proportional to the exposure time), the optimum weight to use for combining the j -th pixel of (suitably aligned) frames i in order to maximise the S/N of sky-limited objects is given by:

$$x'_j = \frac{\sum_i w_{ij} x_{ij}}{\sum_i w_{ij}} \quad w_{ij} = c_{ij} f_i / \sigma_i^2 \quad (10)$$

where σ_i^2 is the average noise variance in frame i , x_{ij} is the flux in pixel j on the i -th frame and x'_j is the combined output flux. The effective exposure time is that of s_o . The output confidence map which is proportional to the inverse of the output noise variance, $\sigma_j'^{-2}$ is therefore given by:

$$c'_j = \frac{(\sum_i c_{ij} f_i / \sigma_i^2)^2}{\sum_i c_{ij} f_i^2 / \sigma_i^2} \quad (11)$$

Special cases of this occur when $f_i = 1$, *e.g.* equal length exposures in stable photometric conditions, or the more general Poisson noise limited case, when $f_i / \sigma_i^2 = 1$, and the special variant of this when $f_i = 1$. These cases are given below, prior to re-normalisation:

$$c'_j = \sum_i c_{ij} / \sigma_i^2 \quad c'_j = \frac{(\sum_i c_{ij})^2}{\sum_i c_{ij} f_i} \quad c'_j = \sum_i c_{ij} \quad (12)$$

3.14 Interleaving

The on-sky detector pixel scale for WFCAM (0.4 arcsec/pixel) undersamples good seeing conditions at Mauna Kea. In order to recover some of this lost resolution, WFCAM observations at a particular pointing may optionally consist of a number of micro-stepped exposures (Casali *et al.* 2001). Microstepping is done by shifting the telescope a precise fractional pixel distance. The most common pattern is a 2×2 microstep sequence where the shifts of the offset exposures (in detector pixels) relative to the first are $(N + 0.5; 0)$, $(N + 0.5; M + 0.5)$, $(0, M + 0.5)$, where N and M are integers. All of the UKIDSS surveys make some use of 2×2 interleaving, while the UKIDSS UDS survey uses 3×3 microstepping and interleaving (Dye *et al.* 2006). Interleaving results in new confidence maps (see section 3.13) for the output image and also, obviously, increases the physical storage size of the image by a factor of 4 or 9.

Algorithmically, 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 (Lauer 1999). Caveats are that interleaving does nothing about bad pixels (in fact bad pixels on an interleaved frame will generally affect more objects), and the PSF often varies on short enough timescales to lead to unusually “spikey” interleaved PSFs that require specialised analysis routines to deal with. A problem encountered occasionally in practice, is that because the location of the interleaves has to be precise (within ~ 0.1 pixels) for the concept to work, the pipeline has to assume the offsets in the headers are correct, rather than deriving them, and that all components of the interleave are present. If for some

reason these conditions are not satisfied, which happens a small but still significant number of times, a bizarre interleave pattern results, which is essentially useless.

Early commissioning data (November 2004) was used to compare the expected offsets from many micro-stepped sequences, with the true offsets calculated from detected object positions, to see whether UKIRT was offsetting with the requisite accuracy. In particular, a series of micro-stepped sequences of an uncrowded bright standard star region in the Z-band were used. Object catalogues for all four detectors from each observation were derived. Then using the first image in each microstep sequence as a reference, we could calculate the (x, y) offsets implied by matching the object positions and compare with the expected offsets as implied in the coordinate information in the headers. These tests indicated that offsetting to the required ~ 0.1 pixel accuracy was in general feasible.

3.15 Dither/jitter stacking

NIR detectors suffer from large numbers of bad pixels, cosmic ray hits and other cosmetic effects. In order to remove these, to lessen residual flatfielding problems, and to avoid saturation by the sky background, the usual practice is to split long integrations into sequences of several shorter, possibly co-added, exposures. Rather than repeating these with each pixel looking at exactly the same sky position, a series of offset exposures are made, usually at integral pixel positions.

This is similar to microstepping, but allows removal of bad pixels and spurions like cosmic-ray events. Combining dithered images with rejection, but without shifts, allows for robust estimates of the sky and fringe patterns (if any) which are useful early on in the reduction process. After removal of the instrumental signatures, dithered images are accurately co-located and combined. Within the pipeline, processing of an observing block, where the offsets are small, involves co-location using internal catalogue generation of (x, y) positions to accurately sub-pixel register (offset) frames prior to stacking; the final image WCS is back-propagated to the individual frame level using these offsets. For external stacking we use the more general purpose derived WCS to drive the co-location operation, with optional use of associated object catalogues for further internal alignment refinement in WCS-transformed pixel space. The first FITS image in the list to be stacked is used as a reference and in WCS-based stacking the other input images are resampled (interpolated)⁷ onto the WCS of the first image.

Prior to pixel rejection, images are scaled (which in general may involve a combination of multiplicative and additive corrections) at the detector level to match the image background in the common overlap region. Pixel rejection is a two-stage process and uses the confidence map (see section 3.13) for each input image to flag the previously known bad pixels and k-sigma clipping (with respect to a sky background calibrated Poisson noise model for the frame, which uses a robust MAD-based sky noise estimator

to define the reference base noise level) to remove cosmic rays and other transient events. All k-sigma clipped pixels are then re-examined to determine if the clipped event was significant with respect to the local *rms* variation in the (preliminary) stacked final image. This latter stage is necessary to avoid clipping out the tops of objects (*e.g.* stars) due to the inevitable pixellation of the flux distribution caused by a combination of subtle seeing variations, steep image gradients and imprecise coordinate alignment.

The surviving pixels are combined using inverse variance weighting, derived from a combination of the input confidence map and the average noise properties of the image, with optional inverse seeing weighting, *i.e.* $\propto \text{seeing}^{-2}$, and weighting according to measures of the throughput (currently based on exposure time but which could make use of ZP information). Since for WFCAM we are in the Poisson noise dominated regime and the exposures in an observing block are taken close together with equal exposure times, the output confidence map is simply the sum of the confidence values of each pixel surviving clipping, renormalised to a median value of 100% (see section 3.13).

3.16 Image mosaicing/tiling

As the focal plane of WFCAM is sparsely populated with detectors spaced $\approx 95\%$ of a detector size apart, for contiguous coverage and analysis of large areas of sky it is necessary to take at least four sets of exposures at large offsets. Tiling, or mosaicing, these to generate a large area image is necessary for detailed two-dimensional study of objects subtending large angles on the sky (*e.g.* nebula in the Galactic plane or nearby external galaxies). However, if contiguous coverage only requires uniform catalogues of (small) detected objects, which is the more general case, then the alternative is to combine catalogues of detected objects from individual “pawprints” rather than produce catalogues from mosaiced images.

Pros and cons of using tiles/mosaics as the basic unit of information are:

pros: tiling makes better use of “fuzzy” dither edges around each of the detectors stacks (pawprints) in a contiguous tile since 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 often a required data product.

cons: tiling generally involves non-linear resampling and hence use of sub-pixel interpolation schemes.⁸ Sub-pixel interpolation affects the noise covariance matrix, complicating later processing stages. Sky and seeing variations make it very difficult to achieve a smooth coherent background on the output image and could also lead to problematic disjoint PSF variations.

⁷ The default interpolation uses nearest-neighbour resampling. Alternatives are to use a Drizzle-like interpolation or higher order interpolating variants such as cubic-splines.

⁸ 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.

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 possibility is to keep the same projection used to describe the telescope+camera, for WFCAM 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 section 4.8.2) which has to be allowed for at the tiling stage.

Since it is much simpler to “tile” detected object catalogues to make unique lists of objects over large areas this is the approach we generally recommend and adopt. To generate a unique catalogue, grouping and selecting between duplicate detections is based on astrometric matching (*e.g.* to within 1 arcsec) and retaining the entry with the lower flux error estimates.

4 CATALOGUE GENERATION

The derived object catalogues are stored in multi-extension FITS files as FITS binary tables, one for each image extension with a dummy primary header unit. Each catalogue header contains a copy of the relevant telescope FITS header content in addition to detector-specific information.

Each detected object has an attached set of descriptors, forming the columns of the binary table and summarising derived position, shape and intensity information. During further processing stages ancillary information such as the sky properties, seeing, average stellar image ellipticity, are derived from the catalogues and stored in the FITS headers attached to each catalogue extension. In addition to being the primary astronomical products, the catalogues and associated derived summary information form the basis for astrometric and photometric calibration and quality control monitoring.⁹

The standard catalogue generation software, (*e.g.* Irwin 1985; Irwin 1996), 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. A subsequent further processing pipeline adds PSF estimation and PSF fits for each object plus a generalised Seric 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). The catalogue generation software makes direct use of the confidence maps for object detection and parameterisation and involves the following series of operations *e.g.* Irwin (1985):

- estimate the local sky background over the field and track any variations at adequate resolution, typically ≈ 30 arcsec to eventually remove them;
- detect objects/blends of objects and keep a list of pixels belonging to each blend for further analysis;
- parameterise the detected objects, ie. perform astrometry, photometry and shape analysis.

4.1 Background analysis

The possibly varying sky background is estimated automatically, prior to object detection, using a combination of robust iteratively clipped estimators.

Any variation in sky level over the frame is dealt with by forming a coarsely sampled background map grid. Within each background grid pixel, typically equal to 64×64 detector pixels (≈ 30 arcsec on sky), an iteratively k-sigma clipped median value of “sky” is computed based on the histogram of flux values within the grid pixel zone. A robust estimate of sigma can be computed using the Median of the Absolute Deviation (MAD) from the median (*e.g.* Hoaglin *et al.* 1983). This is then be further processed to form the frame background map (*e.g.* Irwin 1996).

After removing the varying background component by interpolating onto the image pixel grid, a similar robust estimate of the average sky level and sky noise per pixel is made. This forms part of the quality control measures and also helps to determine the detection threshold for object analysis.

4.2 Object detection

Individual objects are detected using a standard matched filter approach (*e.g.* Irwin 1985, Irwin 1996). Since the only images difficult to locate are those marginally above the sky noise, assuming constant noise is a good approximation (after factoring in the confidence map information) and the majority of these objects will have a shape dominated by the point spread function (PSF), which thereby defines the filter to use.

Potential images are defined in the matched filter map as regions of simply-connected pixels above the detection isophote, relative to the local sky background. On the first pass the sky background is estimated, as outlined previously, and on the second pass the sky level (interpolated if necessary) plus the threshold, define those pixels that require further examination. Spurious images or noise are rejected using a combination of the requirements for a pixel to be above threshold and for it to be connected with enough neighbours to meet a chosen minimum size criterion. The threshold is usually set to be some fixed multiple (≈ 1.5) of pixel-level sky noise above sky. Once a potential image has been located, the original pre-detection filter intensities of the pixels are used for subsequent analysis. This latter requirement ensures that any derived image parameters are unaffected by the blurring action of the detection filter. Of the two most commonly used isophotal connectivity paths,

⁹ The detailed format and content of the object catalogues are described in <http://www.ast.cam.ac.uk/vdfs/docs/catalogues.pdf>

the four nearest neighbour route (N,S,E,W) seems to give better performance than the eight-fold neighbour approach (N,S,E,W + NW,SE,NE,SW). The latter type of connectivity tends to produce more spurious pixels in the outer parts of an image leading to somewhat noisier image parameters. Overlapping objects are automatically deblended by making using a series of isophotes above the detection level (*e.g.* Irwin 1985).

4.3 Object parameterisation

The following image parameters can be computed efficiently and are directly used as part of the image quality control and calibration analysis. In all the following equations, $I(x_i, y_i)$ denotes the intensity of the i th pixel located at position x_i, y_i , and the sum is over all connected pixels.

Isophotal Intensity – the integrated flux within the boundary defined by the threshold level; i.e. the 0th object moment

$$I_{\text{iso}} = \sum_i I(x_i, y_i) \quad (13)$$

For Gaussian images, this is related to the total intensity by the factor $(1 - I_t/I_p)^{-1}$, where I_p is the peak flux and I_t the threshold level (all relative to sky).

Position – computed as an intensity-weighted centre of gravity; i.e. 1st moments

$$\begin{aligned} x_o &= \sum_i x_i \cdot I(x_i, y_i) / \sum_i I(x_i, y_i) \\ y_o &= \sum_i y_i \cdot I(x_i, y_i) / \sum_i I(x_i, y_i) \end{aligned} \quad (14)$$

Covariance Matrix – the triad of intensity-weighted 2nd moments is used to estimate the eccentricity/ellipticity, position angle and intensity-weighted size of an image

$$\begin{aligned} \sigma_{xx} &= \sum_i (x_i - x_o)^2 \cdot I(x_i, y_i) / \sum_i I(x_i, y_i) \\ \sigma_{xy} &= \sum_i (x_i - x_o) \cdot (y_i - y_o) \cdot I(x_i, y_i) / \sum_i I(x_i, y_i) \\ \sigma_{yy} &= \sum_i (y_i - y_o)^2 \cdot I(x_i, y_i) / \sum_i I(x_i, y_i) \end{aligned} \quad (15)$$

The simplest way to derive the ellipse parameters from the 2nd moments is to equate them to an elliptical Gaussian function having the same 2nd moments. It is then straightforward to show (*e.g.* Stobie 1980) that the scale size, $\sqrt{\sigma_{rr}}$, is given by, $\sigma_{rr} = \sigma_{xx} + \sigma_{yy}$; the eccentricity, $ecc = \sqrt{(\sigma_{xx} - \sigma_{yy})^2 + 4 \cdot \sigma_{xy}^2} / \sigma_{rr}$; and the position angle, θ is defined by, $\tan(2\theta) = 2 \cdot \sigma_{xy} / (\sigma_{yy} - \sigma_{xx})$. The ellipticity $e = 1 - b/a$, which is simpler to interpret for estimating potential image distortions (*e.g.* trailing), is related to the eccentricity by $e = 1 - \sqrt{(1 - ecc)/(1 + ecc)}$

Areal Profile – a variant on the radial profile, which measures the area of an image at various intensity levels. Unlike a radial profile, which needs a prior estimate of the image centre, the areal profile provides a single pass estimate of the profile

$$\text{Areal Profile} \rightarrow T + p_1, T + p_2, T + p_3, \dots, T + p_m \quad (16)$$

where $p_j; j = 1, \dots, m$ are intensity levels relative to the threshold, T , usually spaced logarithmically to give even sampling.

The peak height, I_p , is a useful related addition to the areal profile information and is defined as

$$I_p = \max[I(x_i, y_i)]_i \quad (17)$$

or alternatively measured by extrapolation from the areal profile if the image is saturated.

The areal profile provides a direct method to estimate the seeing of objects in an image by enabling the average area of stellar images at 1/2 the peak height, $\langle A \rangle$, to be estimated. The seeing, or FWHM, is then given by $\text{FWHM} = 2 \sqrt{\langle A \rangle / \pi}$.

Finally a series of aperture fluxes are required for object morphological classification (see below)

Aperture flux - we can write this as the integrated flux within some radius r of the object centre

$$I_{ap}(r) = \sum_{i \in r} I_i - N \times \text{sky} \quad (18)$$

where N is the effective number of pixels (which may be non-integral due to the soft-edge apertures used). A series of these is used to define the curve-of-growth, $I_{ap}(r) - v - r$, for each object for later use on both object classification and for deriving aperture corrections for stellar objects. We discuss aperture fluxes further in the next section.

Petrosian flux - is the integrated flux within some radius $k \times r_P$ with $k = 2$ and r_P the Petrosian radius as defined in Yasuda *et al.* (2001).

Kron flux - is the integrated flux within the radius $k \times r_K$ with $k = 2$ and r_K the Kron radius as defined in Bertin & Arnouts (1996).

Hall flux - is the integrated flux within the radius $k \times r_H$ with $k = 5$ and r_H the Hall radius as defined in Hall & Mackay (1984).

4.4 Improved centroiding

An alternative to standard centre-of-gravity centroid estimation that is particularly appropriate for NIR imaging where sky noise dominates virtually all embedded objects, is to use an approximation of full maximum likelihood fitting to generate close to optimal position estimates. It is possible to produce results almost as good as those achieved from full profile fitting, even for overlapping objects, by using an additional weighting term, w_i such that,

$$x_o = \sum_i x_i \cdot w_i \cdot I(x_i, y_i) / \sum_i w_i \cdot I(x_i, y_i) \quad (19)$$

where $w_i = \phi(x_i, y_i)$ is a “model” function of the local surface (a similar equation holds for y_o). Although the “model” is in principle unknown (and usually would need detailed iterative model fitting to generate), we have found that an excellent approximation to it is provided by the pseudo-image produced at the matched detection filter stage (Irwin 1996; Evans, Irwin & Helmer 2002). In this case a single pass operation provides a result almost equivalent to full maximum likelihood analysis of, what may be a complex blend of

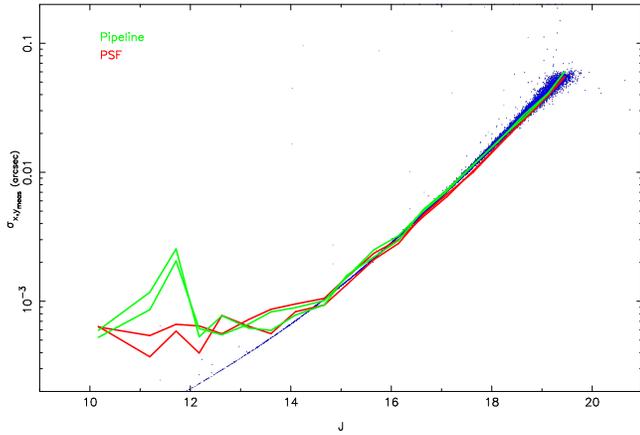


Figure 13. A comparison of PSF and weighted centroid estimation (see text) for a simulated set of WFCAM J-band 2×2 interleaved data taken in stable seeing conditions. Blue points are individual theoretical x,y error estimates, converted to arcsec. The red and green solid lines show median-average binned measured position x,y errors as a function of magnitude for full PSF fitting and the improved centroid estimator. The shot noise at the bright end is due to low number statistics. For real data other factors (*e.g.* differential atmospheric turbulence, detector array irregularities) conspire to produce a lower practical limit of ≈ 5 mas.

objects, and automatically deals well with both stellar and non-stellar images, delivering effectively optimal astrometric precision. We show an example of this in figure 13 comparing full PSF fitting to the above weighted centroiding for a simulated star field.

4.5 Aperture fluxes as a panacea

It is straightforward to show that for the typical range of PSF shapes encountered in practice (*e.g.* some blend of Gaussian, exponential, Moffat), an aperture of radius \approx FWHM delivers almost optimal (80-90%) photometry compared to more detailed (idealised) PSF fitting (*e.g.* Irwin 1996; Naylor 1998).¹⁰

In general this is only the case for isolated images and to work well in more crowded regions several extra “tricks” have to be employed. The first of these is general and makes use of soft-edged apertures, whereby the flux in pixels lying across the aperture boundary is proportioned pro-rata. The second is to use the image detection information for neighbouring regions to flag, and not make use of, those pixels corrupted by external groups of detected objects. The third is to note that, in general, object fluxes from a pixel image blend of overlapping sources are best solved for simultaneously. This is directly equivalent to maximum likelihood or least-squares PSF fitting solely for the fluxes (the object coordinates are already close to optimally determined), where in this case the PSF is a simple top-hat function. With this approach bad pixels and regions of extremely low confidence can also be readily flagged and avoided.

¹⁰ Furthermore, in the limit of bright images dominated by photon noise, rather than sky or readout noise, weighted PSF fitting is equivalent to aperture summation (*e.g.* Irwin 1996).

We find in practice that with these extra refinements there is little to be gained in most practical situations from PSF fitting (see for example figure 19), although somewhat ironically, the main drawback of not using PSF fitting in the standard pipeline is that it is more difficult to generate reliable position error and flux error estimates. In fact we have to make use of theoretical errors based on photon noise statistics, image shape properties and sky estimation errors (Irwin 1985).

One further consideration for using a range of aperture fluxes is paramount. The error in the aperture flux due to inaccurate sky background estimation varies as

$$\delta I_{ap}(sky) = N \times \delta sky = N \times \sigma_{rms} / \sqrt{M} \quad (20)$$

where M is the effective number of pixels used in the sky estimation, and σ_{rms} is the *rms* sky noise per pixel. This of course assumes sky background gradients can be estimated and removed accurately. For faint objects, where sky level errors have most impact, the *rms* error in the aperture flux is mainly caused by sky photon noise and is given by $\delta I_{ap}(rms) = \sqrt{N} \sigma_{rms}$. Therefore requiring that systematic errors in sky estimation contribute significantly less than *rms* photon noise errors implies that $\sqrt{M/N} \gg 1$ *i.e.* the scale size of the background estimator must be \gg than the scale size of the aperture. This effectively rules out locally derived sky background estimation from, for example, an annulus around the image, and favours the more global sky estimation approach that we have adopted.

4.6 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 , $\sqrt{2}r$, $2r$, where r is typically fixed as the median seeing for the site+telescope+camera. For the case of WFCAM we have adopted a set of fixed aperture sizes which use $r = 1$ arcsec as the base unit (up to a maximum size of $12r$) since this also allows comparison with other standard size aperture fluxes.

For classification all detector-level catalogues for each pointing and/or passband are processed independently. Objects are classified based on their overall morphological properties, specifically the curve-of-growth of their flux distribution, and their ellipticity as derived from intensity-weighted second moments. The average stellar locus on each detector in these parameter spaces is generally well-defined and is used as the basis for a null hypothesis stellarness test for use in morphological classification

The classification is primarily based on comparing the curve-of-growth of the flux for each detected object with the well-defined curve-of-growth for the general stellar locus. This latter is a direct measure of the integral of the point spread function out to various radii and is independent of magnitude if the data are properly linearised, and if saturated images are excluded. In using this property the classifier further assumes that the effective PSF for stellar objects is constant over each detector, although individual detectors are allowed to have different PSFs.

The reference stellar loci are defined from the discrete curve-of-growth of the aperture fluxes by analysing the difference in magnitude (or flux ratio) between different pairs of

apertures as a function of magnitude. In practice, the aperture with radius r is used as the fixed reference and also defines the internal magnitude (flux) scale. The linearity of the system implies that the position of the stellar locus for *any function* of the aperture fluxes is independent of magnitude (at least until images saturate). Therefore marginalising the flux ratios over magnitude yields one-dimensional distributions that can be used to greatly simplify automatically locating the stellar locus. With the location fixed, the median of the absolute deviation from the median (MAD) provides a solid measure of the scatter about this locus as a function of magnitude, at least until galaxies dominate in number. This process is repeated iteratively for each distribution, using 3-sigma clipping to remove non-stellar outliers, until satisfactory convergence is reached.

After convergence the equivalent Gaussian sigma is estimated using $\sigma_{gauss} = 1.48MAD$ and by this means each of the image shape descriptors (in this case flux ratios or ellipticity) can be renormalised to follow a zero-median, unit variance Gaussian-like $N(0, 1)$ distribution.

The discrete curve-of-growth of the flux for each object is then compared to that derived from the (self-defining) locus of stellar objects, and combined with information on the ellipticity of each object, to generate the overall detector-level classification statistic. The combination (essentially a weighted sum of the normalised signed distributions) is designed to preserve information on the “sharpness” of the object profile and is finally renormalised, as a function of magnitude, to produce the equivalent of an overall $N(0, 1)$ measure.

In practice measures derived from real images do not exactly follow Gaussian distributions. However, by combining multiple normalised distributions (with well-defined 1st and 2nd moments), the Central Limit Theorem works in our favour such that the resulting overall statistic is Gaussian-like to a reasonable approximation and hence can be used with due care as the likelihood component of a Bayesian Classification scheme, making optional use of prior knowledge.

Objects lying within $2-3\sigma$ of the stellar locus (*i.e.* of zero) are generally flagged as stellar images, those below -3 to -5σ (*i.e.* sharper) as noise-like, and those above $2-3\sigma$ (*i.e.* more diffuse) as non-stellar. An example of the distribution of the classification statistic as a function of (Y-band) magnitude for a UKIDSS Large Area Survey (LAS) pointing is shown in figure 14.

Although the discrete classification scheme is based on the $N(0, 1)$ measure of stellarness it has several overrides built in to attempt to make it more reliable. For example, adjustments to the boundaries at the faint-end (to cope with increased *rms* noise in the statistic) and at the bright-end (to cope with saturation effects) are also made, while the overall image ellipticity provides a further check.

A by-product of the curve-of-growth analysis and the classification is an estimate of the average PSF aperture correction for each detector for those apertures (up to and including $4r$, which includes typically $\sim 99\%$, or more, of the total stellar flux) used in deriving the classification statistic. Accurate assessment of the aperture correction to place the (stellar) fluxes on a total flux scale is a crucial component of the overall calibration. We find that this method of deriving aperture corrections contributes $\leq \pm 1\%$ to the

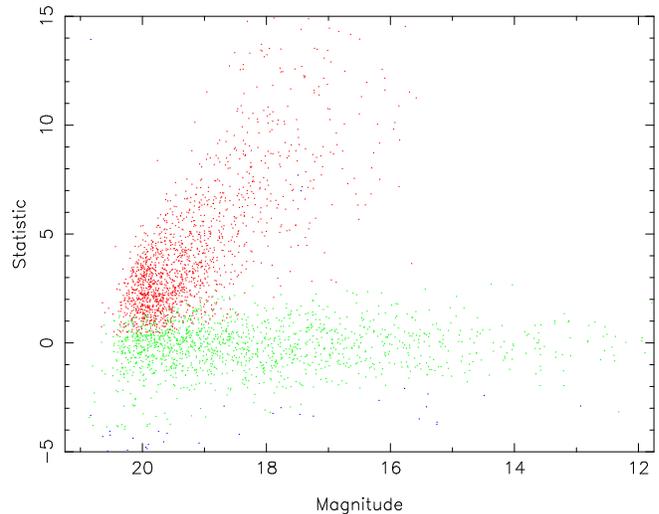


Figure 14. Colour-coded variation of the classification statistic as a function of Y-band magnitude for an LAS pointing; red symbols denote sources classified as non-stellar, blue as noise-like and green as stellar. Although the discrete classification is primarily defined by the statistic, various overrides based on, *e.g.* ellipticity, are also employed to reassign some object classes.

overall photometry error budget and also provides a useful first order seeing correction for non-stellar sources. Further by-products of the morphological classification process are improved estimates of the seeing and average PSF ellipticity from making better use of well-defined stellar-only sources. These parameters are required for quality control monitoring of telescope performance and “atmospheric” seeing.

4.7 Astrometric Calibration

Astrometric calibration is a multi-stage process and aims to provide each image, and any derived catalogues, with a World Coordinate System (WCS) to convert between pixel and celestial coordinates. This happens in the pipeline in two generic stages.

An initial WCS based on knowledge of the instrument, *e.g.* orientation, field-scale, telescope pointing, is embedded in the FITS headers, with telescope-specific information in the primary header and detector-specific information in the secondary headers. This serves to locate each detector image to within a few to several arcsec, depending on the pointing accuracy of the telescope and model parameters. The essential information required 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.

Given a rough WCS for the processed frames, a more accurate WCS can be defined using astrometric standards. We have based our calibration on the 2MASS point source catalog (Skrutskie *et al.* 2006) for several reasons: it is an all-sky NIR survey; it is calibrated on the International Celestial Reference System (ICRS); it provides at least 100 or more suitable standards per pointing; it is a relatively recent epoch (mid-1990s) minimising proper motion problems; the

global systematics are better than 100mas over the entire sky (Zacharias *et al.* 2003); and for 2MASS point sources with signal:to:noise >10:1 the *rms* accuracy per source is $\lesssim 100\text{mas}$.¹¹

From the optical design studies of both WFCAM we know that, to a good approximation, the astrometric distortion is well described by a radially symmetric polynomial distortion model.¹² The effective radial scale r' is related to the true scale r by

$$r' = r + k_3 r^3 + k_5 r^5 + \dots \quad (21)$$

where the k_5 term is usually negligible. If r is given in radians, then the k_3 term has a measured value of ≈ -50 (in units of radian/radian³). This type of distortion is covered by the ZPN projection Calabretta & Greisen (2002) (*i.e.* ARC + polynomial radial distortion). The coefficients for this are encoded in the secondary FITS headers using keywords like PV2.1 and PV2.3.

To refine the astrometric calibration, the 2MASS standard stars and the catalogued objects are first projected into standard coordinates centred on the tangent point of the telescope pointing (*e.g.* Smart 1965). The catalog star projection includes the known WCS distortion model for the system. A binary search over a small grid of offsets centred on the expected location is used to find the best starting point for the refinement with the search radius for a possible match adaptively computed from the 2MASS and object catalogue source densities. Given a list of matched sources the 2MASS standard coordinates are then distorted onto the WFCAM detector system (equation 21) and a linear 6 constant solution is found between pixel coordinates and (distorted) standard coordinates with respect to the optical axis. This is done on a per-detector basis using

$$\xi' = ax' + by' + c \quad (22)$$

$$\eta' = dx' + ey' + f \quad (23)$$

where a, b, d, e encode for scale(s), rotation and shear, c, f are offsets and x', y' denote pixel coordinates shifted to the centre of the matched star distribution for numerical stability in the solution. The refinement is iteratively computed and applied using sequential k-sigma clipping of outlying matched sources until convergence is reached, usually after a few iterations.

Note that the distortion model and the location of the optical axis with respect to each detector is kept fixed and is only updated globally by analysing stacked residuals from many pointing solutions. This decouples the “constant” component of the astrometric solution from the “variable” parts and helps ensure stable least-squares plate solutions. Typical *rms* residuals per star are better than 100mas and are mainly dominated by the 2MASS coordinate accuracy.

Although the *rms* residuals demonstrate that the *rms*

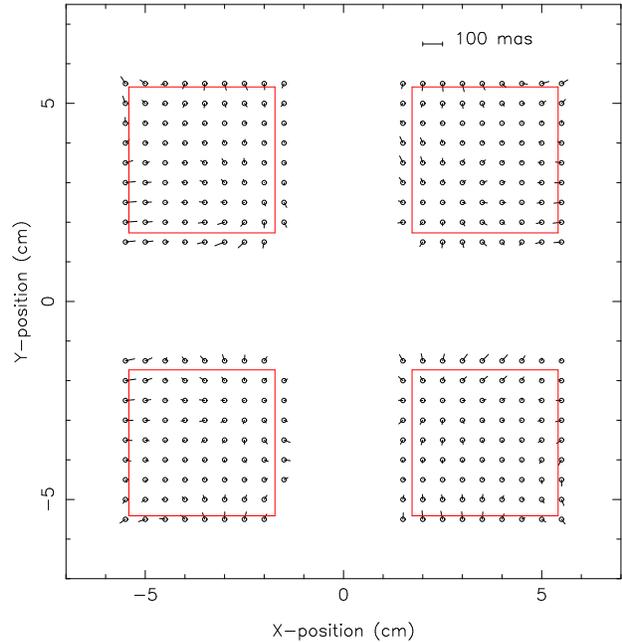


Figure 15. Stacked average astrometric residuals with respect to 2MASS with the idealised positions of the detectors overlaid in red.

accuracy of the plate solution is well within the goal of 100mas, this still leaves possible systematics. We have investigated these by using the overlap between adjacent pointings and also by stacking the 2MASS residuals from sets of many nights of data, taken at large numbers of sky positions to minimise effects of residual systematics in 2MASS. Both of these tests demonstrate that the average systematic residuals over the entire field are significantly better than 100mas. Figure 15 shows an example of the stacked residuals from a 1/2 million K-band objects from the first two weeks of 05B WFCAM observations. There are still residual low level systematics present but the figure, in addition to being a diagnostic of the problem, demonstrates a possible further refinement. The impact of residual systematics could be reduced still further by using a look-up table to interpolate a correction for the repeatable spatially-dependent component. Another possibility would be to incorporate a higher order radial distortion correction term (k_5).

For WFCAM the k_3 coefficient is marginally wavelength-dependent varying by $\approx \pm 10\%$ from the Z-band through to the K-band, and therefore requires a separate k_3 distortion term for each passband used.

4.8 Effects of field distortion

4.8.1 Differential offsets

Although this type of radial distortion generally presents no problem for accurate calibration of individual pointings, it can lead to various complications when stacking data taken at different locations, *e.g.* dither sequences. This is caused by the differential non-linear distortions across individual detectors possibly being comparable to, or larger than, the pixel size of the detector. In these cases stacking involves resampling and interpolation of some form, rather than just

¹¹ Our actual requirement on using a 2MASS star in the astrometry is that it should have a photometric S/N ratio > 10:1 in all of J,H,K. This ensures that the 2MASS astrometry *rms* errors are around 100mas or less for the majority of the point sources used (*e.g.* Irwin 1985). Even in high latitude fields there are still ≈ 50 suitable 2MASS stars per detector to use for the astrometric refinement.

¹² For more information see <http://www.ast.cam.ac.uk/vdfs/>

a simple offset. While resampling is inevitable in combining pointings to form contiguous tiled regions, it may be avoided at earlier stages, such as in stacking individual detector dither sequences, by suitably limiting dither offsets and thereby both simplifying and speeding up the data processing.

The effective local scale for a cubic radial distortion model is given by $dr'/dr = 1 + 3k_3r^2$, which for WFCAM, respectively implies a scale change of $\approx 1\%$ near 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 10.1 arcsec offset at the edge (cf. central pixel scale of 0.40 arcsec). For the outer layer of detectors this may lead to a distortion that is a large fraction of a pixel across each detector during a dither sequence. In turn, this implies that non-linear resampling during stacking may be necessary. To cope with this we are implementing a range of interpolation schemes that offer a trade off between maintaining independent pixel noise and resolution degradation.

4.8.2 Effect of Scale Change on Photometry

One of the contributors to the spatial systematics in the photometry is the effect of the pixel scale change as a function of position. The aim of conventional flatfielding is to create a flat background by normalising out perceived variations from (assumed) uniformly illuminated frames. If the sky area per pixel changes then this is reflected in a systematic error in the derived photometry.

However, since it is often much simpler to deal with “flat” backgrounds, this problem is either usually ignored or corrected during later processing stages, together with other systematic photometry effects. The effect is simplest to envisage by considering what happens to the area of an annulus on sky when projected onto the detector focal plane. From the 3rd order distortion model a sky annulus of $2\pi r ds$ becomes $2\pi r' dr'$ on the detector, which leads to a relative area of

$$(1 + k_3 \times r'^2) \cdot (1 + 3k_3 \times r'^2) \approx (1 + 4k_3 \times r'^2) \quad (24)$$

or in other words roughly 4x the linear scale distortion.

Furthermore, 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 all the photometric systematics in one operation. We discuss in the next section a practical method for achieving this.

4.9 Photometric calibration

We give here a brief overview of the photometric calibration developed for WFCAM (see Hodgkin *et al.* 2006 for more details).

The internal gain-correction, applied at the flatfielding stage, places all the detectors on a common zeropoint (ZP) system (at least to first order *i.e.* ignoring colour equation variations between the detectors). Given a stable instrumental setup, the apparent variation of ZP then directly measures the change in “extinction” without the need to rely on extensive standard field coverage over a range in airmass.

Longer term trends in ZPs due to, for example, accumulation of dust on the optical surfaces can be readily decoupled from shorter term (nightly) variations.

For any given observation of a star in a particular passband

$$m^{cal} = m^{inst} + ZP - \kappa(X - 1) = m^{std} + ce^{std} + \epsilon \quad (25)$$

where ZP is the zeropoint in that passband (ie. the magnitude at airmass unity which gives 1 count/second at the detector), m^{cal} is the calibrated instrumental magnitude, m^{inst} is the measured instrumental magnitude ($-2.5 \times \log_{10}[\text{counts/sec}]$), κ is the extinction coefficient, X is the airmass of the observation, ce^{std} is the colour term to convert to the instrumental system and ϵ is an error term. This assumes that the second-order extinction term and colour-dependency of κ are both negligible. By robustly averaging the ZPs for all the matching stars on the frame an overall ZP for the observation can be obtained.

Photometric calibration is currently based on 2MASS, using linear colour equations to convert from 2MASS to the WFCAM instrumental system. A single 2MASS zero-point (ZP) solution for every catalogued pointing is generated and allows monitoring of effective ZPs at the 1–2% level. The internal gain correction that is applied during flatfielding means that only a single ZP per pointing is required and enables monitoring of the apparent ZP difference between detectors as a contributory diagnostic for non-photometric conditions. The distribution of ZPs in each passband per night is also used to make an estimate of the overall photometric quality of the night and is illustrated in figure 16.

During each night at 1–2 hourly intervals a set of observations centred on UKIRT FS primary calibration stars Hawarden *et al.* (2001), Persson *et al.* (1998) are made to independently monitor the WFCAM 2MASS calibration. Results have been very promising and suggest that the 2MASS-based calibration is delivering frame-by-frame photometric ZPs (with automatically factored-in extinction tracking) at the ± 1 –2% level and even delivers this level of accuracy on the majority of non-photometric nights. That this approach works well is mainly due to the huge effort the 2MASS survey team made in ensuring a reliable all-sky calibration (Skrutskie *et al.* 2006).

A further advantage of being able to use standards from every pointing is that we can use stacked residuals from large series of photometric solutions to directly investigate spatial variations in the photometric calibration. These may be due to a combination of scattered light causing illumination effects in flatfielding, PSF distortions across the focal plane, or astrometric distortion (see section 4.8.2). As an example of this, figure 17 shows the average spatial photometric residuals from a large number of different J-band pointings taken during June 2005 which were all processed using the same twilight flatfields. A correction for astrometric radial distortion is automatically applied during the 2MASS calibration and amounts to $\approx 1\%$ at the outermost corners.

Monitoring the spatial systematics in the photometry is still ongoing but one early conclusion is that it has negligible effect on the overall derived zero-points which are a robust average of the data from all detectors for each calibrated pointing.

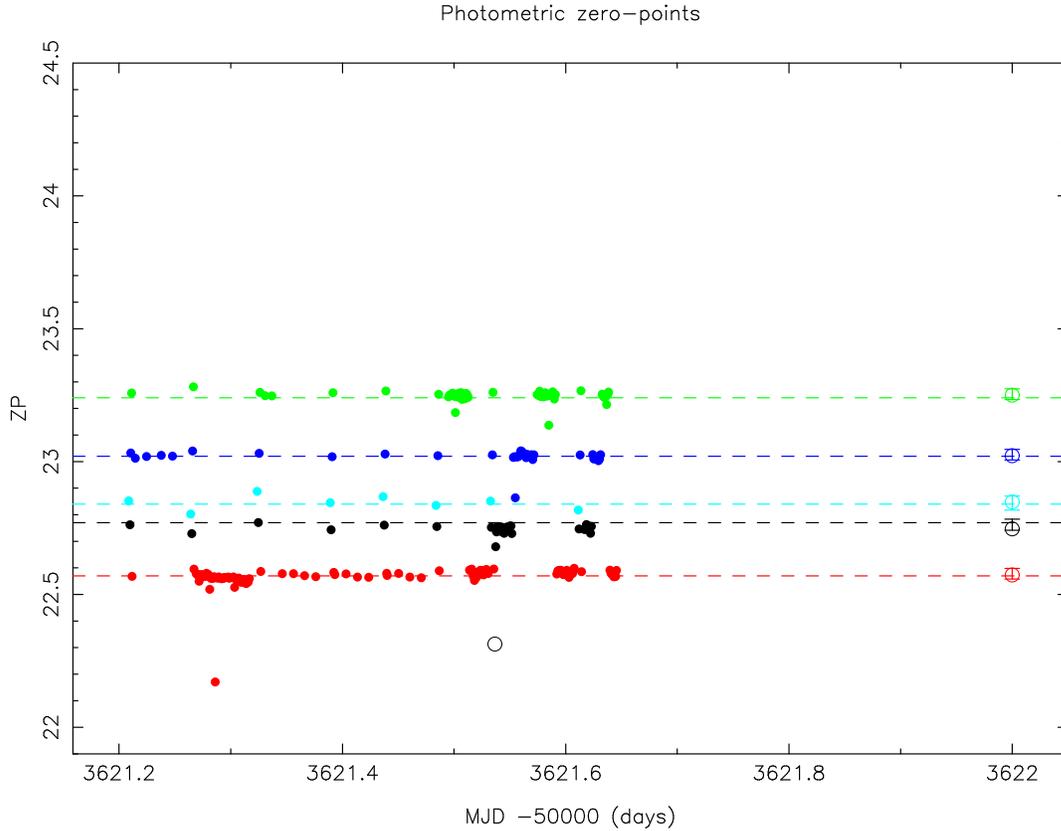


Figure 16. An example of the WFCAM photometric calibration for the night of 20050908. Each filled circle (red K-band, green H-band, blue J-band, cyan Y-band, black Z-band) is a catalog-based individual ZP. Open circles denote either badly trailed or poor seeing data. The dashed lines are nominal ZPs for each passband. Open circles, with error bars, to the right of the figure show the median ZPs for the night with error bars equal to the average robust scatter, *i.e.* a measure of the photometric quality of the night. The outlying red point is from an incomplete observation group.

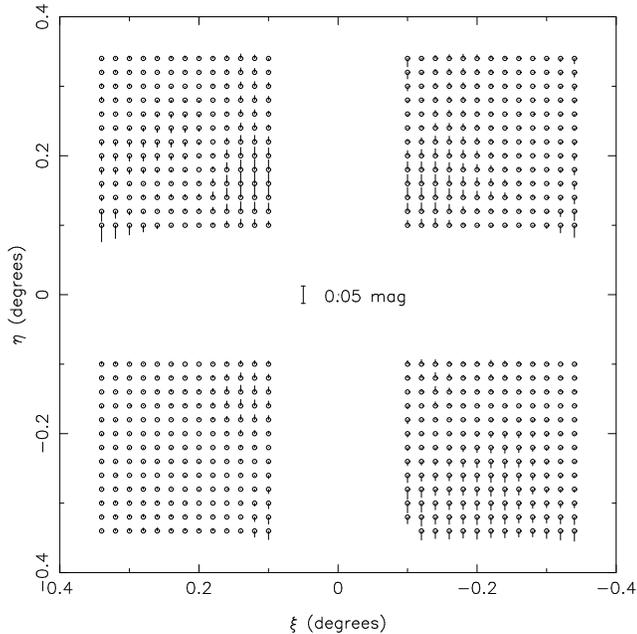


Figure 17. Spatial variation of J-band photometric residuals for June 2005 WFCAM data. The *rms* residual in the binned data is 0.015 mag.

5 PIPELINE OPERATIONS

Having provided a description of the algorithms implemented in the WFCAM pipeline, in this section we give a brief summary of how the pipeline is operated to reduce a typical night of WFCAM data.

There are five main stages, the first and the last require user interaction, while the central three are fully automated.

(i) **Preparation:** Raw data are copied and uncompressed on-the-fly to the working directory, calibrations (flatfields, confidence maps) are retrieved from the calibration frame library. Master darks for each exposure time/read mode combination are created and, after that, the single dark frames, focus frames etc. are removed from the working directory. A preview script is run on the directory producing a log describing what reduction steps will be carried out for each observing block and what calibration will be used. The log file permits a straightforward identification of incomplete groups (aborted observing blocks), missing calibrations and/or sky frames (see item 3), which can be optionally fixed if feasible. Additional information to drive the pipeline (*e.g.* how to construct sky correction frames, which sky correction frame to use for a given science frame) is written to an instruction file.

(ii) **Group processing 1:** During this stage the pipeline groups images according to the observing block they be-

long to. For each observing block all the images are dark-subtracted, flatfielded, decurtained and cross-talk corrected. Note that all the images in any given group will have been observed with the same filter.

(iii) **Sky correction:** The aim of this stage is to form the sky frames (see section 3.10). All the science frames are re-grouped by filter and a number of sky frames are created. Given the short time-scale of sky variation, in general, there will be several sky frames for each filter and exposure time. The number of sky frames per filter can be tuned via three pipeline parameters: minimum number of images to form a sky frame (usually 4), maximum number of images to form a sky frame (8-12) and minimum angular separation between the images used to form a sky frame (generally 0.3 arcmin, but can be reduced if necessary). All the images are then sky subtracted (after proper re-scaling of the sky frame) using the most appropriate sky frame.

(iv) **Group processing 2:** After having sky-corrected all the images, the pipeline reverts to observation group processing to carry out the last stages of the reduction: interleaving (if needed), stacking, catalogue generation, WCS fitting and photometric calibration.

(v) **Data quality check:** After pipeline processing is over, data quality control (DQC) plots and information are created from the object catalogues (see section 6). Eventually the data is compressed and moved to the final reduced data storage location and flagged as ready for transfer by the WSA group in Edinburgh.

6 QUALITY CONTROL

Every processed WFCAM data frame has 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.

Other quality control measures include night-time observing conditions and the status of the telescope + instrument, detectors/filters/controllers etc.. This information is included in the FITS headers at the summit as part of the observing process.

The summit pipeline, and later the standard processing pipeline, derives further quality control measures directly from the data and records them for later use. These measures include: sky brightness, sky noise, average stellar ellipticity, average FWHM seeing measure, stellar aperture corrections, number of spurious images, astrometric errors, WCS parameters and first pass photometric calibration; and on a nightly basis, derived local and global photometricity flags, limiting magnitude estimates, nightly zero points and extinction measures.

All the data quality information for every data product and input science frame are stored in a Data Quality Control (DQC) database. The data is ingested in the database straight after the processing integrity has been checked. The ingestion procedure is also used to perform a number of additional checks on the FITS header contents.

A web front-end to the database enables further CASU internal data quality checks by allowing: easy searches for

data at particular pointings, observing dates and/or with multiple constraints on the DQC parameters. The database front-end is further complemented by an image cut-out service that is used to create on-the-fly object postage stamps and full-chip previews. The QC database also provides the means to monitor the progress of pipeline processing and the progress of the UKIDSS surveys.¹³

7 EXAMPLE SCIENCE DATA

In this section we illustrate the quality of the pipeline output by showing two science examples which highlight the diversity of astronomical images that the pipeline is being used to automatically analyse. Both examples come from UKIDSS survey data taken during 05A, the first semester of science operations of WFCAM.

In both examples shown the following steps were taken to select and combine the pipeline catalogue output.

(i) Only frames passing sensible quality control criteria were used, *i.e.* seeing <1.0 arcsec, average stellar ellipticity < 0.2, reasonably good photometric night from inspection of the nightly photometry quality control information.

(ii) The pipeline generated object catalogues were matched detector by detector using a 2.5 arcsec search radius with an iterative 6 constant linear solution applied to matched objects, in (rectilinear) standard coordinate space, to improve the differential astrometry. Typical differential astrometry averages ≈ 50 mas for all the regions used.

(iii) Objects matching within 1.0 arcsec were considered reliable matches, *i.e.* the probability of a spurious mismatch is low even in crowded regions. Objects matching in the range 1.0 - 2.5 arcsec were flagged as unreliable using the classification index as a peg.

(iv) Colour-magnitude and colour-colour diagrams were constructed from all objects classified as stellar, or probably stellar, in all bands. In addition, for the colour-colour diagrams, objects were only plotted if their estimated magnitude errors were <0.1 mag in each passband.

(v) For tiles and larger areas, a unique object catalog was constructed by searching for duplicate entries within 1 arcsec and retaining the entry with the lowest magnitude error.

The results for the $13.7' \times 13.7'$ region from a single detector from the GPS survey of the M17 region (figure 19) show that even in crowded regions with complex background variations the pipeline photometry gives good results. While as a contrast, the LAS tile results shown in figure 20 illustrate not only the interesting inter-play between the locus of most stars and galaxies in NIR passband parameter spaces, but also verify that searching for rare objects (*e.g.* high red-shift quasars, or L- and T-dwarfs) is a feasible proposition.

8 SUMMARY

TBD.

¹³ <http://casu.ast.cam.ac.uk/surveys-projects/wfcam/data-processing/>.



Figure 18. A $10' \times 10'$ region in M17 from a K-band image taken from the UKIDSS GPS. Note in particular the large variation in reddening across even this small region.

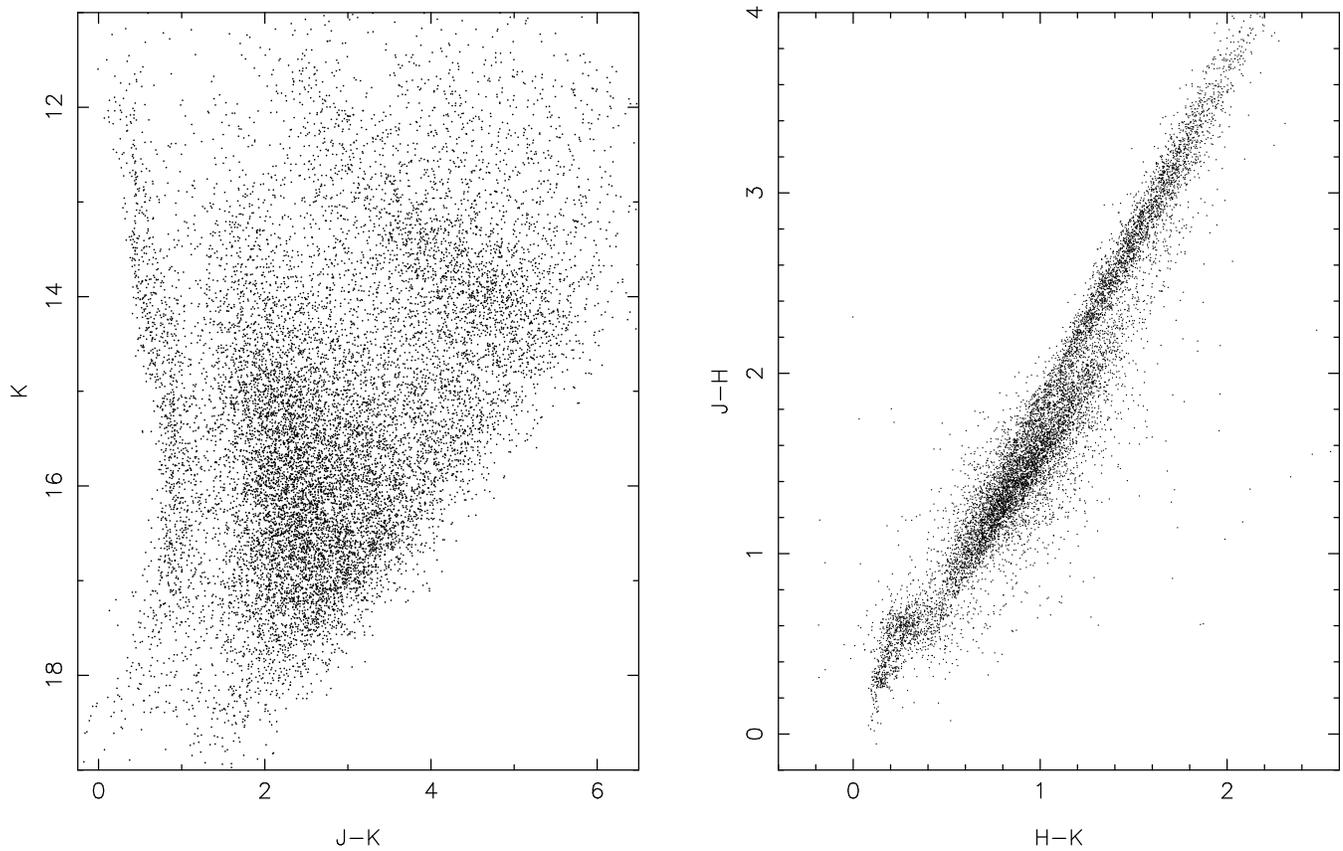


Figure 19. Colour-magnitude and colour-colour diagrams for M17 from a single WFCAM detector including the region illustrated above. Even in crowded complex regions like these the pipeline produces reliable photometry.

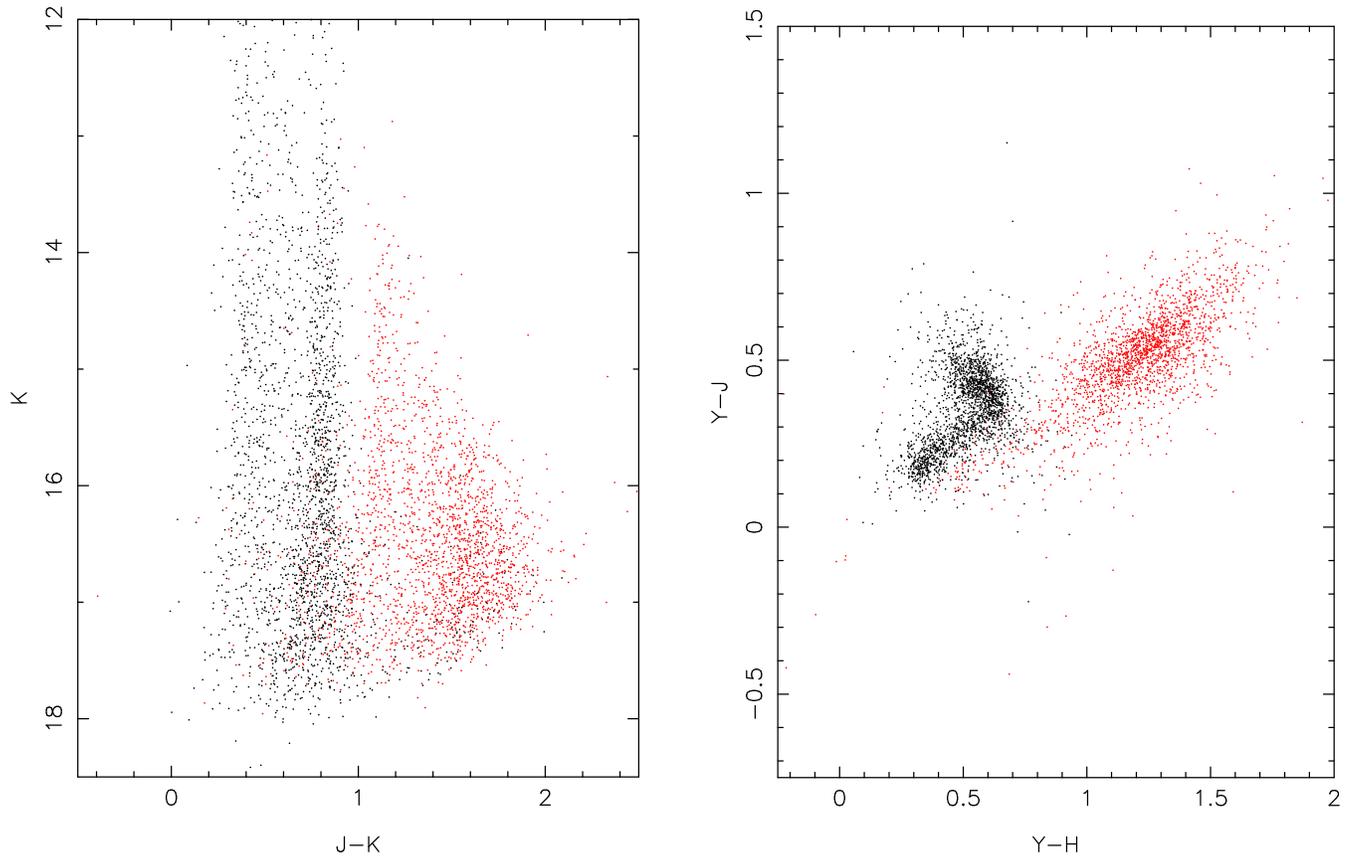


Figure 20. Photometry for a complete LAS tile of ≈ 0.8 sq deg in Y,J,H,K bands. Stellar-classified sources are shown as black dots, non-stellar as red illustrating the quality of the morphological classification and also the ability of NIR photometry to help discriminate between galactic stars and background galaxies.

ACKNOWLEDGMENTS

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

REFERENCES

- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
 Calabretta, M. R., & Greisen, E. W. 2002, *A&A*, 395, 1077
 Casali, M., et al. 2001, *ASP Conf. Ser.* 232: *The New Era of Wide Field Astronomy*, 232, 357
 Casali, M., et al. 2006, *MNRAS*, submitted
 Dye, S. et al. 2006, *MNRAS*, in preparation
 Emerson, J. P. 2004 *SPIE*, 5493, 411
 Evans, D. W., Irwin, M. J., & Helmer, L. 2002 *A&A*, 395, 347
 Finger, G. 2006 *SPIE*, 6276, in press
 Greisen, E. W., & Calabretta, M. R. 2002, *A&A*, 395, 1061
 Hall, P., & Mackay, C. D. 1984, *MNRAS*, 210, 979
 Hambly, N. C. 2004 *SPIE*, 5493, 423
 Hambly, N. C., et al. 2006, *MNRAS*, in preparation
 Hawarden, T. G. 2001, *MNRAS*, 325, 563
 Hoaglin, D. C., Mosteller, F., & Tukey, J. W. 1983, *Wiley Series in Probability and Mathematical Statistics*, New York: Wiley, 1983, edited by Hoaglin, David C.; Mosteller, Frederick; Tukey, John W.,
 Hodgkin, S. T. et al. 2006, *MNRAS*, in preparation
 Irwin, M. J. 1985, *MNRAS*, 214, 575
 Irwin, M. J. 1996, in *7th Canary Islands Winter School*, Ed. Espinosa, J.M.
 Irwin, M. & Lewis, J. 2001, *NewAR*, 45, 105
 Irwin, M. J. 2004 *SPIE*, 5493, 401
 Lauer, T. R. 1999, *PASP*, 111, 227
 Lawrence, A. et al. 2006, *MNRAS*, in preparation
 Naylor, T. 1998, *MNRAS*, 296, 339
 Pence, W. D. 2002, *Proc. SPIE*, 4847, 444
 Persson, S. E. et al. 1998, *AJ*, 116, 2475
 Sabbey, C. N. 1998, *PASP*, 110, 1067
 Sabbey, C. N. 2000, *PASP*, 112, 867
 Skrutskie, M.F., et al. 2006, *AJ*, 131, 1163
 Smart, W. M., 1965, *Spherical Astronomy*, Cambridge University Press
 Stobie, R. S. 1980, *JBIS*, 33, 323
 Yasuda, N., et al. 2001, *AJ*, 122, 1104
 Zacharias, N., et al. 2003, *IAUJD*, 16, 43

APPENDIX A: NON-LINEARITY CORRECTION FOR CDS DATA

There are two general options for non-linearity correction. The first is to compute and apply the non-linearity correction on individual readouts, preferably in the detector

controller system, for data volume considerations, prior to computing the reset-corrected frames. The second method forgoes direct access to the individual frames and instead computes and applies the non-linearity correction at the front-end of the pipeline.

Each option has certain advantages and disadvantages. The former method 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.

The only clear drawback of correcting post-facto is that it would probably be very difficult, if not impossible, to correct multiple NDR gradient fitting output that way. In the following we describe a simple scheme for directly dealing with non-linearity in CDS data.

A1 Correcting for non-linearity

In default CDS reset-read-read (RRR) mode, downstream of the data acquisition system (DAS) the output that we see is

$$\Delta I' = I'_2 - I'_1 = f(I_2) - f(I_1) \quad (\text{A1})$$

where I'_1 and I'_2 denote the non-linear first (*i.e.* the reset-frame) and second readouts respectively and I_1 and I_2 the desired linear quantities. The non-linear function $f(I)$ maps the distortion of the desired linear counts to the non-linear system I' . If we define the inverse transform $g(I')$ that maps measured counts I' to linearised counts I as the inverse operator $g() = f^{-1}()$ then

$$I = g(I') \quad \text{and} \quad I_1 = g(I'_1) \quad I_2 = g(I'_2) \quad (\text{A2})$$

If I'_1 and I'_2 were directly available this is a one-to-one mapping and can be done efficiently and accurately using Look Up Tables (LUTs). This is the conventional way of implementing the correction prior to other image manipulation operations.

However, if I'_1 and I'_2 are not separately available and all we have to work from is the difference $\Delta I'$ then a simple LUT transformation is not possible.

For example, taking the simplest case where the illumination level across the detector has not changed during the course of the RRR and no on-board co-addition is happening then, in principle given only ΔI and knowledge of the timing of the RRR operations we can deduce I_1 and I_2 by using the effective integration time for each to estimate their scaling to the measured difference ΔI such that,

$$I_1 = k\Delta I \quad \text{and} \quad I_2 = (1+k)\Delta I \quad (\text{A3})$$

Unfortunately, the ratio k will not be constant for the non-linear quantities I'_1 and I'_2 forcing us to adopt a scheme along the following lines.

Given $\Delta I'$ and defining the non-linear operator $f()$ as a polynomial with coefficients a_m (typically up to quartic) we have

$$\Delta I' = \sum_m a_m (I_2^m - I_1^m) = \sum_m a_m [(1+k)^m \Delta I^m - k^m \Delta I^m] \quad (\text{A4})$$

The quantity we want ΔI is buried in the non-linearity of the RHS and we have to solve an equation like this for every pixel. This is possible, and relatively simple to program using something like a Gauss-Seidel iterative scheme, but is more inefficient than a direct mapping.¹⁴

If we wanted to use a completely general LUT approach we would require a 2D LUT for all possible values of I_1 and I_2 *i.e.* 65k × 65k in size, or 4.3×2 Gbytes. Most likely we would need a different correction for each “channel” making a total of 128 (WFCAM) × 8.6 Gbytes = 1.1 Tbytes of LUT for the cameras! Of course if the range of values of k is limited via exposure time quantisation this decreases the size of the total no. of LUTs required considerably for the constant illumination case, but would be an ugly and possibly impractical solution.

Practical considerations (*e.g.* data volume), suggest two alternative solutions for non-linearity correction: either correct the individual frames directly in the DAS by measuring and downloading the appropriate LUTs, or polynomial coefficients, to the DAS; or use a non-linear inversion on the reset-corrected frames as outlined here. This methodology is not generally applicable, *eg.* to multi-NDR/gradient fitting readouts, but is directly applicable to coadded (or coaveraged) frames of the same exposure times, assuming constant illumination over the series.

For reset-corrected data, the non-linear inversion is competitive with complex operations on LUTs and much simpler to implement. It also has the added advantage of removing all aspects of the non-linearity correction from the DAS. The main disadvantages are the method is restricted to CDS RRR mode, and if the illumination level is rapidly varying (*eg.* twilight) the effective scale factors k_i may be hard to compute accurately - although for all realistic practical situations the knock-on effect is likely to be negligible.

A2 Measuring the non-linearity

If all that is available are reset-corrected data from say a time series of dome flats, it is still feasible to directly compute the non-linearity coefficients.

Given a series of measurements $\{i\}$ of $\Delta I'_i$ and using the previous notation and polynomial model

$$\Delta I'_i = \sum_m a_m (I_2^m - I_1^m) = \sum_m a_m \Delta I_i^m [(1+k_i)^m - k_i^m] \quad (\text{A5})$$

where k_i are the exposure ratios under the constant illumination.

In general $\Delta I_i = s t_i$ where t_i is the exposure time of the i th reset-corrected frame in the series and s is a fixed (for the series) unknown scale factor. The k_i are computable from a knowledge of the exposure times and the reset-read overhead, t_i and $\Delta I'_i$ are measured quantities leaving the polynomial coefficients a_m and the scaling s to be determined.

¹⁴ A simulation on a twin processor 3GHz Xeon PC indicates, for typical values, a time of $\approx 1/8$ th second per 2k × 2k detector.

Thus the model is defined by

$$\Delta I'_i = \sum_m a_m (I_2^m - I_1^m) = \sum_m a_m s^m t_i^m [(1+k_i)^m - k_i^m] \quad (\text{A6})$$

and can be readily solved by standard linear least-squares methods using the following sleight-of-hand. Since the scaling s and hence the polynomial solution a_m are coupled, by simply (and logically) requiring in the final solution $a_1 = 1$, computation of s can be completely avoided.

Rewriting the previous equation in the following form makes this more apparent

$$\Delta I'_i = \sum_m (a_m s^m) t_i^m [(1+k_i)^m - k_i^m] = \sum_m b_m t_i^m [(1+k_i)^m - k_i^m] \quad (\text{A7})$$

where now b_m are the coefficients to be solved for. The final step is to note that

$$a_m = b_m / s^m = b_m / b_1^m \quad (\text{A8})$$

since by definition $a_1 = 1$.

A3 Individual frames available

If both reset and data frames are available then the problem is much simpler since it can be phrased in the following way

$$I_1 = \sum_m a_m I_1'^m = s t_1 \quad (\text{A9})$$

$$I_2 = \sum_m a_m I_2'^m = s t_2 \quad (\text{A10})$$

where t_1 and t_2 are the known effective exposure times. Given a series of measurements $\{i\}$ these can either be solved as a coupled pair of linear least squares minimisations or by forming

$$\Delta t = \sum_m a_m / s (I_2'^m - I_1'^m) \quad (\text{A11})$$

as a standard linear least-squares problem. Again the requirement that $a_1 = 1$ automatically defines s . These coefficients can then be directly applied to the observed flux to give the desired linearised units.