GaBoDS: The Garching-Bonn Deep Survey

IV. Methods for the Image reduction of multi-chip Cameras demonstrated on data from the ESO Wide-Field Imager*

T. Erben[1], M. Schirmer[1,2], J. P. Dietrich[1], O. Cordes[1], L. Haberzettl[1,3], M. Hetterscheidt[1], H. Hildebrandt[1], O. Schmithuesen[1,3], P. Schneider[1], P. Simon[1], E. Deul[4], R. N. Hook[5], N. Kaiser[6], M. Radovich[7], C. Benoist[8], M. Nonino[9], L. F. Olsen[10,8], I. Prandoni[11], R. Wichmann[12], S. Zaggia[9], D. Bomans[3], R. J. Dettmar[3], and J. M. Miralles[13,1]

[1]Institut für Astrophysik und Extraterrestrische Forschung (IAEF), Universität Bonn, Auf dem Hügel 71, D-53121 Bonn, Germany
[2]Isaac Newton Group of Telescopes, Apartado de correos 321, 38700 Santa Cruz de La Palma, Tenerife, Spain
[3]Astronomisches Institut der Ruhr-Universität-Bochum, Universitätsstr. 150, D-44780 Bochum, Germany
[4]Leiden Observatory, Postbus 9513, NL-2300 RA Leiden, The Netherlands
[5]Space Telescope European Coordinating Facility, European Southern Observatory, Karl-Schwarzschild-Strasse 2, D-85748 Garching bei München, Germany
[6]Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA
[7]INAF, Osservatorio Astronomico di Capodimonte, via Moiariello 16, 80131 Napoli, Italy
[8]Laboratoire Cassiopée, CNRS, Observatoire de la Côte d'Azur, BP4229, 06304 Nice Cedex 4, France
[9]INAF, Osservatorio Astronomico di Trieste, Via G. Tiepolo 11, 34131 Trieste, Italy
[10]Copenhagen University Observatory, Copenhagen University, Juliane Maries Vej 30, 2100 Copenhagen, Denmark
[11]Instituto di Radioastronomia, CNR, Via Gobetti 101, 40129, Bologna, Italy
[12]Hamburger Sternwarte, University of Hamburg, Gojenbergsweg 112, 21029 Hamburg, Germany
[13]European Southern Observatory, Karl-Schwarzschild-Strasse 2, D-85748 Garching bei München, Germany

Received June 15, 2004; accepted June 16, 2004

Abstract. We present our image processing system for the reduction of optical imaging data from multi-chip cameras. In the framework of the Garching Bonn Deep Survey (GaBoDS; Schirmer et al. 2003) consisting of about 20 square degrees of high-quality data from WFI@MPG/ESO 2.2m, our group developed an imaging pipeline for the homogeneous and efficient processing of this large data set. Having weak gravitational lensing as the main science driver, our algorithms are optimised to produce deep co-added mosaics from individual exposures obtained from empty field observations. However, the modular design of our pipeline allows an easy adaption to different scientific applications. Our system has already been ported to a large variety of optical instruments and its products have been used in various scientific contexts. In this paper we give a thorough description of the algorithms used and a careful evaluation of the accuracies reached. This concerns the removal of the instrumental signature, the astrometric alignment, photometric calibration and the characterisation of final co-added mosaics. In addition we give a more general overview on the image reduction process and comment on observing strategies where they have significant influence on the data quality.

Key words: methods: data analysis - techniques: image processing

1. Introduction

During the last decade, optical Wide-Field Imaging has become one of the most important tools in observational

Correspondence to: T. Erben: terben@astro.uni-bonn.de

* Based on observations made with ESO Telescopes at the La Silla Observatory

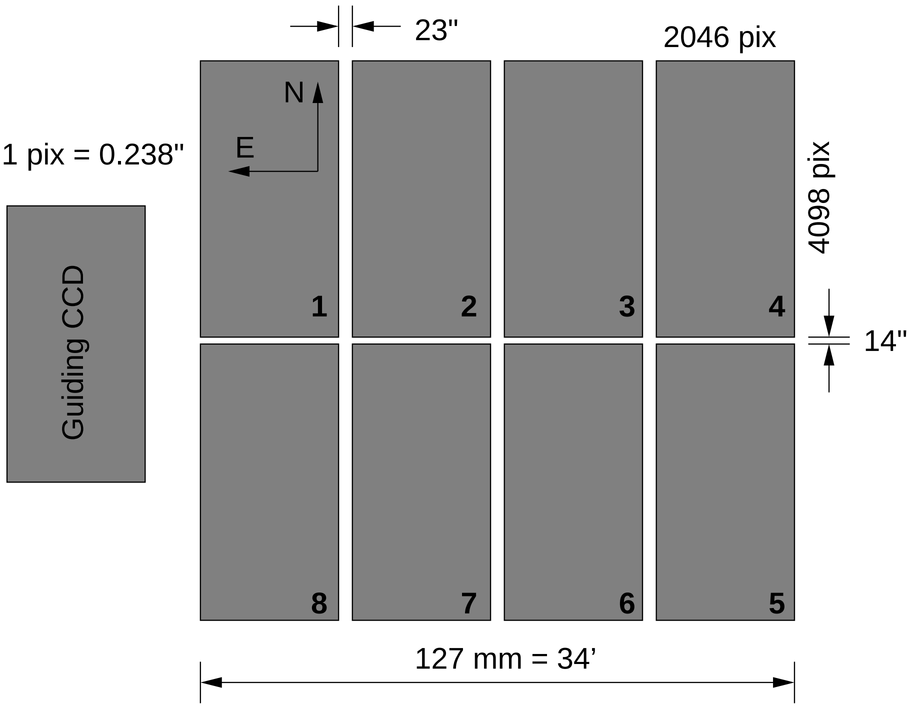

Fig. 1. Shown is the CCD layout of WFI@MPG/ESO 2.2m (not to scale). The mosaic design of CCDs to obtain a larger field-of-view generally confronts the user with new issues compared to the treatment of single-chip cameras. The possible overlap of sources lying on different detectors on dithered exposures significantly complicates the astrometric and photometric calibration and the subsequent co-addition process. First, the large field-of-view leads to notable optical distortions from the field centre to the edges (see also Fig. 11). These distortions have to be known accurately for each CCD to create a correctly stacked mosaic. Second, for homogeneous photometry we need to take into account different quantum efficiencies and hence different photometric zero points of individual chips. Finally, dithering leads to non-uniform noise properties in the co-added image (due to the gaps, sensitivity and gain differences; see also Fig. 26). The knowledge of the relative noise properties between individual pixels in the co-added mosaic is important in the subsequent source catalogue creation process (see Fig. 27).

astronomy. The advances in this field are closely linked to the development of highly sensitive detectors based on charge coupled devices (CCDs). These detectors have grown in size from a few hundred pixels on a side to the currently largest arrays of about 4000×4000 pixels. The need for ever larger fields-of-view and deeper images, and the technical constraints in the manufacturing processes of larger CCDs led to the development of multi-chip cameras in the mid 90's. Hereafter, we refer to them simply as *Wide Field Imagers*, or WFIs. The production of such a mosaic detector array with well aligned CCDs in three dimensions is very difficult. The misalignments of individual chips, and the multi-chip nature by itself, are the main reasons that the data reduction process is significantly more complicated compared to the treatment of single-chip cameras. Although a multi-chip CCD camera seems to be a simple collection of N single-chip cameras at first glance (see Fig. 1) it is, in general, not possible to treat all images belonging to a certain detector independently during the complete reduction process.[1] To obtain an accurately aligned co-added image on the whole WFI scale we have to properly account for the mosaic structure with its gaps, not perfectly aligned chips, and a large field-of-view that needs to be treated as a curved surface instead of a flat plane. We have to apply sophisticated techniques going considerably beyond the work that has to be done for a single-chip camera with a small field-of-view. In addition, a homogeneous photometric calibration and the characterisation of the noise properties in a co-added mosaic poses new challenges on calibration techniques. Aside from the technical issues related to the mosaic structure of new cameras, the increasing quality of new optical instruments with their high sensitivity and throughput demand the use of state-of-the-art algorithms to optimally exploit the obtained data scientifically.

Up to now, no *standard procedure* to process and treat WFI data is established among observers. The main reason is that WFI capabilities have only been available to a broader community for about half a decade and we are still in the process of fully understanding the properties of the new instruments and developing the necessary tools for the data handling, and also for analysing reduced images. Another issue is the high data

[1] This is still possible for certain scientific objectives not needing any image co-addition (as variability studies in single frames for instance) or if data are obtained in a very compact dither pattern so that different chips never overlap. As we will show in the course of the paper, compact dither patterns have in general severe drawbacks on the final image quality.

flow and the demands on computer hardware connected already to small and medium-sized WFI projects. Only for about two years high-end, and for the occasional observer affordable, desktop PCs are equipped adequately.

That these issues are non-trivial is illustrated by the relatively small number of publications based on WFI projects compared to the amount of data acquired. For example, after the WFI at the 2.2m MPG/ESO telescope (hereafter WFI@2.2m) became operational in January 1999, only 35 refereed papers based on its data appeared until the end of 2002. Meanwhile, the rate has risen[2], but is still significantly behind those of other ESO instruments which are under similar pressure by proposers as WFI@2.2m.

In this paper we present the methods and algorithms we use to process WFI data from the ESO Wide-Field Imager. The presentation is organised as follows: In Sect. 2 we introduce the basic concepts of the GaBoDS pipeline and the philosophy behind the choices we made. We discuss the advantages and the disadvantages that arise thereof. The pre-reduction process (i.e. the removal of the instrumental signature from the data), which is in principle identical for single- and multi-chip cameras, is described in Sect. 4. Details of the astrometric and photometric calibration of the data are presented in Sect. 5. An explanation of our adopted scheme to deal with the inhomogeneous noise properties in co-added mosaic data is given in Sect. 6, followed by the image co-addition methods in Sect. 7. We perform quality control checks on co-added data in Sect. 8 and draw our conclusions in Sect. 9.

No astronomical pipeline will produce the best possible result with data obtained in arbitrary conditions and strategies. Hence, besides a pure description of algorithms, our presentation also contains guidelines how WFI data should be obtained to achieve good results.

We assume familiarity with data processing of optical imaging data. In this publication we focus on describing the algorithms used for all necessary processing steps. Where the reduction differs significantly from well established algorithms for single-chip cameras (this mostly concerns the astrometric alignment and the photometric calibration) we give a scientific evaluation and a thorough estimation of the accuracies reached.

We note that we mainly work on data from WFI@2.2m and most of the examples and figures in this paper refer to data from this instrument. The reader has to be aware that the quoted results, the accuracies and the overall usability of our algorithms can vary significantly when being applied to data sets from other cameras. At critical points we come back to this issue in the text.

[2] The *ESO Telescope Bibliography* turns up a total of 130 refereed papers until April 2005.

2. Pipeline characteristics

2.1. Scientific motivation

While all optical data need the same treatment to remove the instrumental signature (bias correction, flatfielding, fringe correction etc; see Sect. 4), the subsequent treatment of the images strongly depends on the scientific objectives and on the kind of data at hand. Our primary scientific interests are weak gravitational lensing studies (see e.g. Bartelmann & Schneider 2001 for an extended review) in deep blank-field surveys. These studies mainly depend on shape measurements of faint galaxies. To ensure that the light distribution is deformed as little as possible by the Earth's atmosphere and the optics of the telescope, weak lensing data are typically obtained under superb seeing conditions at telescopes with state-of-the-art optical equipment. For a proper treatment of those data, the main requirements on a data processing pipeline are the following:

- We need to align very accurately galaxy images of subsequent exposures that are finally stacked. This involves an accurate mapping of possible distortions (see Fig. 11).
- We need the highest possible resolution, and hence a co-addition on a sub-pixel basis (see Sect. 7). Together with the previous step, this is crucial for an accurate measurement of object shapes in the subsequent lensing analysis.
- We need to accurately map the noise properties in our finally co-added images (see Fig. 27). This knowledge enables us to use as many faint galaxies as possible and to push the object detection limit. This also requires that the sky-background in the co-added image is as flat as possible.

The algorithms we use are chosen to go from raw images to a final co-added mosaic with the objectives described above. The responsibility of the pipeline ends with the co-addition step. We are fully aware that the chosen procedures and algorithms may not be optimal for projects having different scientific objectives such as accurate photometry, the investigation of crowded fields or the study of large scale, low surface brightness objects for instance. However, the modular design of our pipeline (see below) allows an easy exchange of algorithms or the integration of new methods necessary for different applications. The main characteristics of our current pipeline are summarised in the following:

- Ability to process exposures from a multi-chip camera with N CCDs. To date we have successfully processed data from: WFI@2.2m, CFH12K@CFHT, FORS1/2@VLT, WFI@AAO, MOSAIC-I/II@CTIO/KPNO, SUPRIME-CAM@Subaru, WFC@INT, WFC@WHT, and several single-chip cameras (e.g. BUSCA@2.2m Calar Alto).
- Ability to handle mosaic data taken with arbitrary dither patterns.

- Precise image alignments without prior knowledge about distortions. The astrometric calibration is performed with the data itself.
- Absolute astrometric calibration of the final images to sky coordinates.
- An appropriate co-addition of data obtained under different photometric conditions.
- Image defects on all exposures are identified and marked before the co-addition process.
- Creation of weight maps taking into account different noise/gain properties and image defects for the co-addition.
- Accurate co-addition on the sub-pixel level/ability to rescale data to an "arbitrary" scale (ability to easily combine data from different cameras/telescopes).

We are currently extending our algorithms to near-IR cameras which will be described in a forthcoming publication (Schirmer et al. in prep.).

2.2. Implementation details

As mentioned above, many tools for WFI data processing are currently under active development and different groups have already released excellent software packages for specific tasks. Hence, we built our own efforts on already publicly available software modules wherever possible. Many of the algorithms used are very similar to those developed for the EIS (ESO Imaging Survey) Wide survey in 1997-1999 (see Nonino et al. 1999). The main pillars of our pipeline are the following software modules:

- **The LDAC software[3]:** The LDAC (Leiden Data Analysis Centre) software was the backbone of the first EIS pipeline. It provides a binary catalogue format (in the form of system-independent binary FITS tables) including a large amount of tools for their handling. Moreover, this module contains software for the astrometric and photometric calibration of mosaic data.
- **The EIS Drizzle[4]:** A specific version of the IRAF package `drizzle` (Fruchter & Hook 2002) was developed for EIS. It directly uses the astrometric calibration provided by the LDAC astrom part and performs a weighted linear co-addition of the imaging data.
- **TERAPIX software[5]:** (Bertin et al. 2002) `SExtractor` is used to obtain object catalogues for the astrometric calibration. Moreover it produces a cosmic ray mask in connection with `EYE`[6] in addition to smoothed and sky-subtracted images at different parts of the pipeline. `SWarp`, the TERAPIX software module for resampling and co-adding FITS images, is used alternatively to `EIS Drizzle` for the final image co-addition.
- **FLIPS[7]:** (Magnier & Cuillandre 2004) `FLIPS` is one of the modules for data pre-reduction. It is optimised to perform operations on large format CCDs with minimal memory requirements at the cost of I/O performance.
- **Eclipse and qfits tools[8]:** (Devillard 2001) We use several stand-alone FITS header tools from the `Eclipse` package to update/query our image header. Moreover, tools based on the `qfits` library offer an alternative to `FLIPS` for image pre-reduction.
- **Astrometrix[9]:** `Astrometrix` (developed at TERAPIX) is another module for obtaining astrometric calibration.
- **IMCAT utilities[10]:** From the `IMCAT` tools we extensively use the image calculator `ic`.

These tools have been adapted for our purposes if necessary and wrapped by UNIX/bash shell scripts to form our pipeline. Our main effort is to provide the necessary interfaces for the communication between the individual modules and to add instrument and science specific modules for our purposes (as a thorough quantification of PSF properties for instance). With our implementation approach we can make use of already well-tested and maintained software packages. The description above also shows that we can easily exchange modules as soon as new algorithms or better implementations for a certain task become available. To further enhance the modularity of the pipeline, we split up the whole reduction process into small tasks (a complete reduction process for WFI@2.2m data typically involves the call to 20-30 different scripts but *superscripts* collecting several tasks can easily be written). This ensures a very high flexibility of the system and that potential users can easily adapt it to their needs or to specific instruments.

Nearly the whole system is implemented in ANSI C and UNIX bash shell scripts (`EIS Drizzle` is written in FORTRAN77 and embedded into IRAF, `Astrometrix` is developed under `Perl+PDL` and parts of our photometric calibration module are implemented in `Python`). Thus we have full control over the source codes and can port the pipeline to different UNIX flavours. Table 1 lists the systems on which we successfully used it so far.

The main disadvantage of building up a pipeline from many different software modules instead of developing a homogeneous system from scratch is that it becomes very difficult to automatically control the data flow or to construct a sophisticated error handling and data integrity checking system. Furthermore, our pipeline so far offers only very limited possibilities to store the history

[3] available at ftp://ftp.strw.leidenuniv.nl/pub/ldac/software/

[4] available via ESO's SCISOFT CD (see http://www.eso.org/scisoft/)

[5] available at http://terapix.iap.fr/soft/

[6] `EYE` (Enhance Your Extraction) allows the user to generate image filters for the detection of small-scale features in astronomical images by machine learning (neural networks). These filters are loaded into `SExtractor` that performs the actual detection. We use such filters to detect cosmic rays in our images.

[7] see http://www.cfht.hawaii.edu/~jcc/Flips/flips.html

[8] available at http://www.eso.org/projects/aot/eclipse/

[9] available at http://www.na.astro.it/~radovich/wifix.htm

[10] available at http://www.ifa.hawaii.edu/~kaiser/imcat/

Table 1. Listed are the different UNIX architectures on which we used our pipeline so far. Besides the system we list the compilers and the co-addition module used. We compared the final co-added images from different systems with respect to flux and shape measurements from extracted objects. The individual reductions do not show significant differences.

System	compiler	co-addition
Linux/i386	gcc-2.95.3	`SWarp`, `EIS Drizzle`
Linux/alpha	gcc-2.96	`SWarp`
Linux/AMD64	gcc-3.3.4	`SWarp`
SunOS	gcc-2.95.2	`SWarp`, `EIS Drizzle`
IBM AIX	gcc-2.95.2	`SWarp`
HP-UX	gcc-2.7.2.2	`SWarp`
OSF1	gcc-2.7-97r1	`SWarp`
IRIX64	gcc-3.0.4	`SWarp`

of the reduction process or to administrate raw and processed image products. Also, formal speed estimates for the throughput of our processing system are low compared to homogeneous systems (see Table 2). Thus, the usability for large, long-term projects of the system is limited at this stage.

With the compactness, the flexibility and the usability of our system for the occasional user[11] we regard our tools complementary to the survey systems developed by other groups such as TERAPIX[12] or ASTRO-WISE[13]. All software modules of our pipeline, that are based on publicly available code, can be obtained by request to the authors. We note that the EIS team recently released a WFI pipeline (Vandame 2002, 2004a, 2004b) with a functionality similar to ours[14]. An alternative, and widely used reduction package for CCD mosaics is the `mscred` environment within `IRAF`. It is described in Valdes (2002).

2.3. Parallel processing

A great deal of the reduction processes is performed independently on individual chips or on a set of images belonging to the same CCD. Only very few steps require information or imaging data from the complete mosaic. This mainly concerns the estimation of a global astrometric and photometric solution for the entire mosaic (which is done on catalogues) and the final image co-addition. Hence, most processing tasks can naturally be performed in parallel if the mosaic data is split up on a detector basis and if a multi-processor architecture is at hand. Our pipeline is designed to perform this kind of parallelisation on a single multi-processor machine as well as on machines connected over a network (such as a Beowulf Cluster for instance). This parallelisation scheme of performing the same task simultaneously on different machine processors rather than using explicit parallel code for the individual tasks has two main advantages:

– No new software has to be developed for the parallel processing. We developed a *parallel manager* that launches and surveys script execution on the individual cluster nodes. The parallel manager is also used in standard, single-node processing so that the pipeline operation is unified in single- and multi-processor mode.
– As the operations on different chips are completely independent and no communication is necessary between the tasks, no message passing/synchronisation or data exchange schemes have to be implemented. We only synchronise the processing after every processing step, waiting until a certain job is done on all chips before the next task is launched.

The major disadvantage in our scheme not allowing a parallelisation in the processing of a given chip is that resources are not used optimally if more nodes/processors than chips are available (nodes are simply staying idle for the complete process) or if the number of nodes is not an integer multiple of chips (some nodes have to work on more chips than others, basically doubling processing time). Nevertheless, we think that the advantages heavily outweigh this disadvantage for our purposes. Moreover, most multi-chip cameras so far have 2^N chips matching typical architectures of computer clusters. In Table 2 we compare processing times from single- and multi-node processing of an observing run.

3. Terminology

We use the following terminology throughout the paper:

– CCD, chip – one of the detectors in a WFI
– exposure – a single WFI shot through the telescope
– image – the part of the exposure that belongs to a particular CCD
– frame – this term is used as a synonym for an image.
– BIAS – an exposure/image with zero seconds exposure time
– DARK – an exposure/image with non-zero exposure time, keeping the shutter closed
– FLAT – an exposure/image of a uniformly illuminated area; this can be the telescope dome giving DOMEFLATs or sky observations during evening and/or morning twilight providing SKYFLATs
– SCIENCE – an exposure/image of the actual target, not a calibration image
– SUPERFLAT – properly stacked SCIENCE data to extract large scale illuminations or fringe patterns.
– STANDARD – an exposure/image of photometric standard stars
– other terms written in CAPITAL letters denote additional images or calibration frames. Their meaning will be clear within the context.

[11] The usage of our pipeline can easily be learned within a few days by users having good experience in the reduction of single-chip cameras, as was experienced by several guests who made use of our visitor programme.
[12] http://terapix.iap.fr/
[13] http://www.astro-wise.org/
[14] http://www.eso.org/eis/

Table 2. Processing times for a WFI@2.2m set consisting of 10 Bias, 12 Dark, 12 Skyflat and 11 Science frames. The third column lists the percentage of the processing time with respect to the single-node case. The number in parentheses gives the percentage if the whole process could be parallelised, i.e. 100% divided by the number of nodes. The fourth column lists the formal number of Science frame pixels processed in one second. The times include the full processing from raw data as they arrive from the ESO archive to the final co-added images [it does not include the time for eyeballing and masking of images (see Sect. 4.10)]. We note that the parts that cannot be parallelised (an estimate for the global astrometric alignment with `Astrometrix`, the global photometric solution and the final co-addition with `SWarp`) and have to be done on a single processor add up to 24 min in our example. This becomes a very significant fraction of the total processing time in the 4 and 8 node cases. Each node is equipped with an Athlon XP 2800+ processor and 1 GB of main memory. For a comparison, for the current EIS pipeline a throughput of 0.4-0.5 Mpix/s on a double-processor machine is quoted (Vandame 2002).

used nodes	time [min]	time [%]	Kpix/s
1	240	100 (100)	50
2	131	54.0 (50)	90
4	77	32.0 (25)	152
8	52	21.7 (12.5)	226

- mosaic – this term is used as a synonym for exposure or for the final co-added image.
- dithering – offsetting the telescope between the exposures
- overlap – images from different CCDs and exposures overlap if the dithering between the exposures was large enough
- stack – a set of n images belonging to the same chip that have been combined pixel by pixel.
- Names of software packages are written in the `TypeWriter` font.

4. Pre-reduction (*Run* processing)

In the following we describe our algorithms for the pre-reduction of optical data, i.e. the removal of the instrumental signature. The first issue is the compilation of data for this step. For most of the effects to be corrected for (instrument bias, bad CCD pixels, flat-field properties) we can safely assume stability of the instrument and the CCD characteristics over several days or even a few weeks. Hence, in many cases we can collect data from a complete *observing run* and benefit from having many images for the necessary corrections from which most are of statistical nature. As described below, the matter can become more complex when dealing with strong fringes in red passbands. Here, the time scales from which SCIENCE data can safely be combined in the process is much shorter, sometimes only a couple of hours. The issue is further discussed in Sect. 4.9. In any case we say that we perform the pre-reduction process on a *Run basis* regardless of how long this period actually is. The pre-reduction is done independently on each CCD of a mosaic camera. Only in one step, the sky-background equalisation (see Sect. 4.6), the action to be performed on a CCD depends on properties from the rest of the mosaic. Hence, unless stated otherwise, each step described below has to be performed on a detector basis. For this part of the pipeline we can use two different software packages, one based on `FLIPS`, the other on `Eclipse`. `FLIPS` is very I/O intensive but has very low memory requirements, while `Eclipse` reduces the necessary I/O to a minimum and operates very efficiently on imaging data by keeping it in virtual memory. Depending on the size of the data set and the computer equipment at hand, one or the other is preferable. In the following we focus on the description of the `Eclipse` package and we mention `FLIPS` where its functionality differs significantly from `Eclipse`.

4.1. Handling the data and the FITS headers

The variety of ways in which header information and raw data from WFIs are stored in FITS files are as large as the number of instruments, i.e. so far there is no established standard on the FITS format of CCD mosaics. In order to cope with the different formats and to unify the treatment, we perform the following tasks on raw data:

1. If the data are stored in Multiple Extension Fits (MEF) files, we extract all the images from them. All subsequent pipeline processing is done on individual images and also our pipeline parallelisation is based on the simultaneous execution of the same task on different images.
2. We substitute the original image headers from the chips by a new one containing only a minimum set of keywords necessary for a successful pipeline processing (see App. A). Especially the astrometric calibration depends on correct header entries.
3. If necessary we flip and/or rotate individual CCDs to bring all images of the mosaic to the same orientation with respect to the sky. Only rotations by 90 degrees, that do not require pixel resampling, are performed.

All these tasks are performed by a `qfits`-based utility.

4.2. Modes, medians, and the stacking of images

In most pre-reduction steps we have to perform robust statistical estimates on a set of pixels following Poisson statistics. At some points we have to estimate the sky-background and its variance in SCIENCE data. A typical distribution of pixel values on a SCIENCE image is shown in Fig. 2. The second common operation is to statistically stack several input images to produce a noise-reduced calibration frame. This master image is then used for subsequent calculations on SCIENCE images. In the following we define several *operators* acting on images:

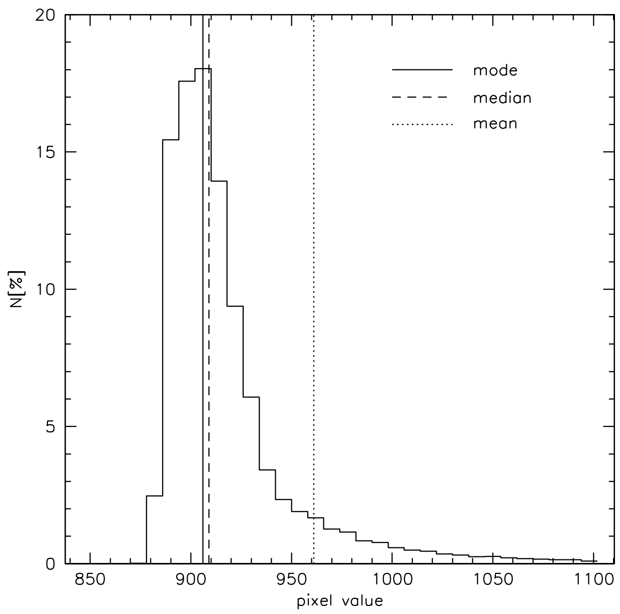

Fig. 2. Shown is a typical pixel value distribution in an astronomical image. The clear peak denotes the brightness of the night sky. Due to the presence of objects the distribution is strongly skewed towards high values. The most accurate estimate for the sky comes from analysing the histogram and determining the peak (the mode of the distribution) directly. The median (for a set of N points the median element is defined so that there are as many greater as smaller values than the median in the set) also gives a very robust (and for our purposes sufficient though biased) estimate for it. Due to the long tail of the distribution, a straight mean is completely useless as an estimate for the sky value. As the values to the left of the mode form half of a Gaussian distribution, an estimate for the sky variance can be obtained from that part.

- *value=median(image):*
  The median of an image is estimated. This is done by collecting a representative sample from the CCDs pixels (we take about 5% of the detector pixels in a region around the CCD centre), sorting the obtained pixel array and returning its middle element (in case the array has an even number of elements we return the mean of the two middle elements).
- *value=mode(image):*
  The mode of an image is estimated. We consider the same representative pixel sub-sample as for the median estimation, build a smoothed histogram and return the peak value.
- *image=rescale(images):*
  This operator rescales a set of images so that they have the same median after the process. The resulting median is chosen to be the mean of the medians (we write *meanmed* for it) from the input images. Hence on each input image we perform the operation *image → image * meanmed / median(image)*. The operator is usually applied to a set of images before they are stacked (as SKYFLATs or SCIENCE images). The median of the stacked image is then by construction equal to *meanmed*.
- *image=stack(images):*
  A set of input images is stacked to produce a noise-reduced *master image*. The following procedure is performed independently on each pixel position. We collect all pixel values from the input images into an array, sort it and reject several low and high values (typically we reject the three lowest and three highest values if we have 15 input images or more). From the rest we estimate the median that goes into the master image. Here `FLIPS` uses a more sophisticated algorithm. On the remaining array it first performs an iterative sigma clipping. It estimates mean and sigma, rejects low and high values (typically pixels lying more than 3.5 sigma below and above the mean) and repeats the procedure until no more pixels are rejected. From the rest it returns the median for the master image.

We will use additional pseudo operators in the following whose meaning and behaviour will be clear within the context. All calculations with these operators are written in *slanted* notation.

4.3. A first quality check, Overscan correction, master BIASes and DARKs

Before any exposure enters the reduction process, we estimate its mode. If this estimate lies outside predefined boundaries, the exposure will not be processed. This rejects most of the completely faulty images (such as saturated SKYFLATs) at the very beginning. After this initial quality check, the first step in the reduction process of each image is the correction of an overall estimate for the BIAS value by considering pixels in not illuminated parts of each CCD (the overscan region). This first-order BIAS correction is done by collecting for each line all pixels in the overscan region, rejecting the lowest and highest values (usually the overscan regions contains about 40 columns and we reject the 5 lowest and 5 highest values) and estimating a straight mean from the rest. This mean is subtracted from the corresponding line. After this correction, the overscan regions are trimmed off the images. For correcting spatially varying BIAS patterns in the FLAT or SCIENCE images, a *master BIAS* is created for each CCD from several individual BIAS exposures:

1. Each BIAS exposure is overscan corrected and trimmed.
2. The master BIAS is formed by *masBIAS=stack(BIAS frames)*

In the same way, a master DARK (*masDARK*) frame is created. So far, we do not correct FLAT or SCIENCE frames for a possible dark current but we use the masDARK for the identification of bad pixels, rows and columns (see Sect. 6).

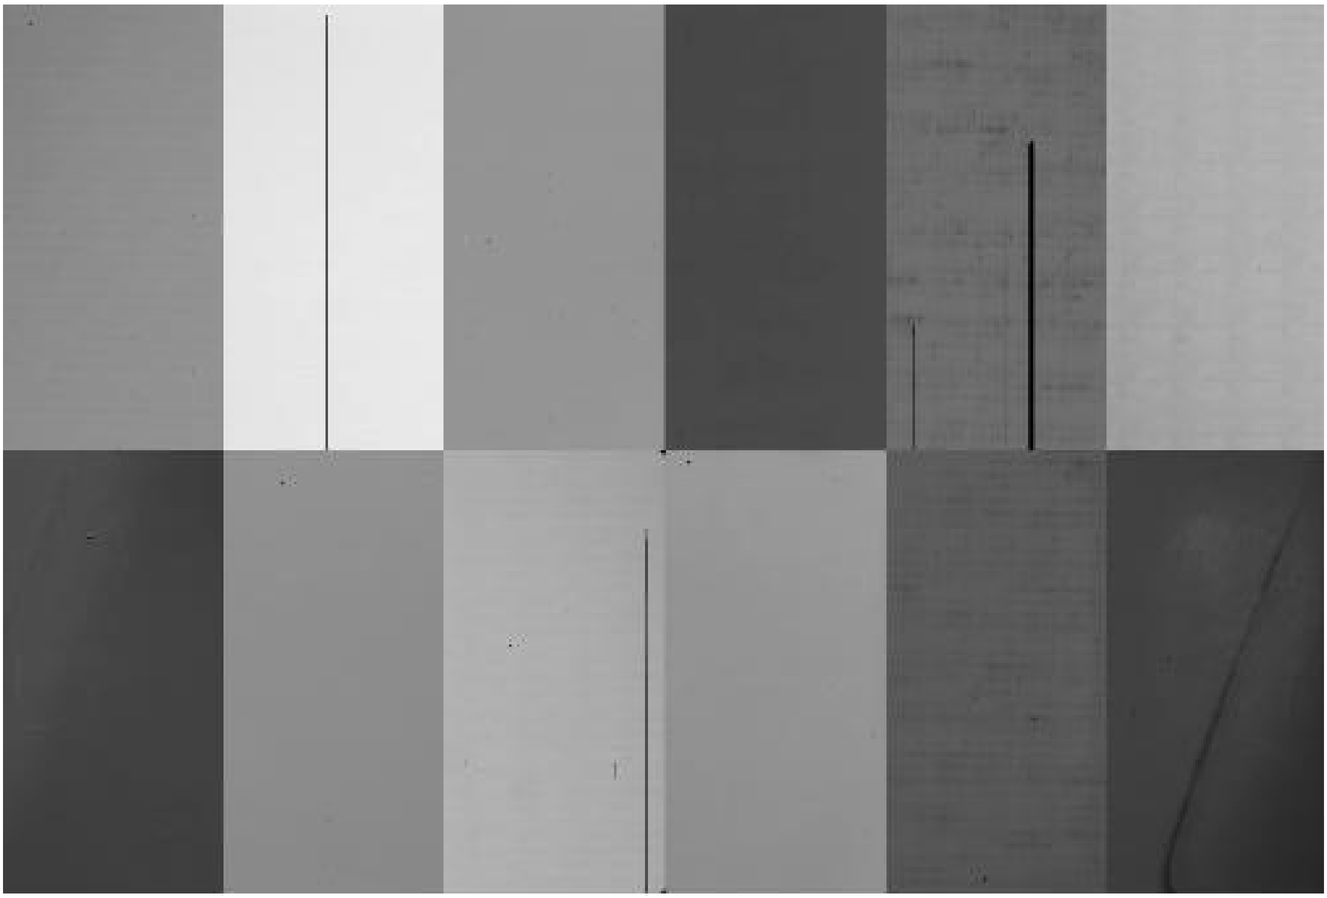

Fig. 3. A skyflat from the CFH12K@CFHT camera (Cuillandre et al. 2000). We note the significant intrinsic sensitivity and gain differences between the chips. The counts of the second left detector in the top row are about a factor 1.3 higher than that of the two chips at the corners in the bottom row. Within a CCD, large-scale gradients are up to 8%. Especially the four chips on the right also show small-scale sensitivity variations across the whole CCD.

4.4. Methods for flat-fielding

The illumination of a CCD and its sensitivity is in general not homogeneous over the detector area. Illumination usually varies on large scales, but sensitivity can change significantly from pixel to pixel (see Fig. 3). In order to overcome these effects and to achieve a homogeneous surface brightness, SCIENCE images have to be flat-fielded. The flat-field pattern typically varies with the observation configuration and hence calibration images have to be obtained for each filter to be used. There are three major possibilities to obtain an estimate for the flat-field-pattern:

1. **DOMEFLATs:**
   One can obtain flat images by pointing the telescope to an illuminated, blank surface in the dome. These DOMEFLATs can be taken during day time, the images do not contain any astronomical sources, we can take as many exposures as we wish and we can adjust the count rate in the frames. This ensures FLATs with a very high signal-to-noise. However, especially for WFIs, the uniform illumination of a large enough area is very difficult to achieve. In addition, lamps emit light at a maximum temperature of around 3000 K so that the spectral distribution is very red. In red filters this can lead to an excess amount of fringing as compared to SKYFLATs. For these reasons, the flat calibration of SCIENCE frames with DOMEFLATs is usually inferior to the two methods described below (especially at the correction of large-scale illumination variations). We only use them if SKYFLATs are not available.
2. **SKYFLATs:**
   One can also obtain flat images by observing the sky with short exposure times during the evening and morning twilight. These flats resemble the night sky during SCIENCE observations better than dome flats, although for WFIs covering an area of $30' \times 30'$ or more we expect variations of the twilight sky over their field-of-view. At many telescopes, an automatic setup to obtain exposures with a desired count-rate are available and hence these flats also have high S/N. See also Tyson & Gal (1993) for a guide to obtain twilight flats if an automatic procedure is not available. However, only a very limited number can be obtained during the twilight phase. This holds especially if several filters have to be calibrated and/or if the readout time of the detector is rather long. Using an exposure time too short in bright twilight can lead to a non-uniform illumination due to shutter ef-

fects; see Sect. 4.9. However, CCDs are usually very stable over a few days or weeks so that a larger number of SKYFLATs taken during several nights can be combined into good master flat images.

3. **SUPERFLATs:**
In addition, one can try to extract the flat-field pattern from the SCIENCE observations itself. If a sufficient number of SCIENCE exposures is at hand (usually more than a dozen), and if the dither pattern was significantly larger than the largest object in the field, each pixel of the camera will see the sky-background several times. Hence, a proper combination of these exposures yields a master FLAT that closely represents the night sky during observations. We refer to such a FLAT as a SUPERFLAT. A straightforward application of this method is often hampered by several factors: During phases of grey and bright time, the night sky shows large gradients and variable reflections in the dome and telescope can occur. Thus, a careful selection of images that go into the SUPERFLAT has to be done. In medium/narrow band filters and in the ultra violet, the counts and hence the S/N of SUPERFLATs in these bands are typically low. Furthermore, the technique is very difficult to apply in programmes where the imaged target has a size comparable to the field-of-view. In this case also large dither patterns cannot assure sufficient sky coverage on the complete mosaic.

As the superflat technique cannot be applied in general we use the following two-stage flat-fielding process:

1. In a first step, the SCIENCE observations are flat-fielded with SKYFLATs (if those are not at hand we use DOMEFLATs instead). This typically corrects the small-scale sensitivity variations and leaves large-scale gradients around the 3% level on the scale of a chip (see Fig. 5).
2. If the data permit it, a SUPERFLAT is created out of the flat-fielded SCIENCE observations, smoothed with a large kernel to extract remaining large-scale variations and is applied to the individual images. For our empty field observations at the WFI@2.2m this leaves typical large-scale variations around the 1% level in U, B, V, R broad band observations (in I the presence of strong fringing often leads to significantly higher residuals).

This two-stage flat-fielding process is very similar to the method adopted by Alcalá et al. (2002).

4.5. The creation of master DOME-/SKYFLATs

The master FLAT (SKYFLAT or DOMEFLAT) for each CCD is created as follows:

1. All individual FLAT exposures are overscan-corrected and trimmed.
2. The masBIAS is subtracted from all images.
3. The FLATs are rescaled to a common median: *rescFLAT=rescale(FLAT)*.
4. We form the master FLAT by *masFLAT=stack(rescFLAT frames)*

4.6. Sky-background equalisation

Within a mosaic, CCDs have different quantum efficiencies, varying intrinsic gains and hence different photometric zero points (see Fig. 3). For the later photometric calibration it is desirable to equalise the photometric zero points in all detectors, which is achieved by scaling all CCDs to the same sky-background.We rescale all chips to the median of the CCD with the highest countrate during flat-fielding. If possible we perform this step within the SUPERFLAT correction as its median estimation is more robust than with the SKYFLATs. This is because SKYFLATs show larger variations in brightness than the SUPERFLATs which are calculated from already flat-fielded data. We estimate that photometric zero points of the mosaic agree with an rms scatter of about 0.01-0.03 mag after sky-background equalisation (see Sect. 5.6).

4.7. SCIENCE image processing, the creation and application of the SUPERFLAT

The individual steps in the SCIENCE frame processing are:

1. The images are overscan-corrected and trimmed. Afterwards the masBIAS is subtracted and the frames are divided by masFLAT. We write *flatSCIENCE=(SCIENCE−masBIAS)/masFLAT*. In the case where no SUPERFLAT correction is applied to the data, masFLAT is rescaled at this step to equalise the sky-background between individual detectors (see Sect. 4.6).
2. For the creation of the SUPERFLAT we first remove astronomical sources from the flatSCIENCE images. To this end we run `SExtractor` (Bertin & Arnouts 1996) on them and create a new set of images where pixels belonging to objects (i.e. above certain detection thresholds) are flagged (*objSCIENCE=flatSCIENCE−OBJECTS*). See also Fig. 4. The flagged pixels are not taken into account in the subsequent processing.
3. The objSCIENCE images are rescaled to a common median *rescobjSCIENCE=rescale(objSCIENCE)*.
4. We calculate the SUPERFLAT [*SUPERFLAT=stack(rescobjSCIENCE frames)*]. If all input pixels of a given position are flagged (i.e. all the images had an object at that position) we assign the *meanmed* value from the objSCIENCE frames to the SUPERFLAT.
5. The SUPERFLAT is heavily smoothed by creating a `SExtractor` BACKGROUND check-image with a background mesh of 512 pixels (see Bertin 2003 on the SExtractor sky-background estimation). This image, the illumination correction image, forms the basis for removing large-scale flat-field variations (*ILLUMINATION=smooth(SUPERFLAT)*).

6. The flatSCIENCE images are divided by the ILLUMINATION image which has been rescaled to equalise the sky-background of the different detectors (see Sect. 4.6). We write *illumSCIENCE=flatSCIENCE/rescaled(ILLUMINATION)*.

For blue passbands not showing any fringing, the pre-reduction ends here and the illumSCIENCE images are used in the subsequent processing. See Fig. 5 for an example of a pre-reduced V-band exposure.

4.8. Fringing correction in red passbands

Fringing is observed as an additional, additive instrumental signature in red passbands. It is most prominent on cameras that use thinned CCDs and hence are optimised for observations in blue passbands. Fringes show up as spatially quickly varying, wave-like structures on the CCDs (see Fig. 6). The *geometry* of these patterns usually does not change with time since this interference effect is created in the substrate of the CCD itself. WFI@2.2m shows fringes with an amplitude of about 1% as compared to the sky-background in broad-band R. In the I- and Z-bands, fringing becomes much stronger and reaches up to about 10% of the night sky. Unfortunately, contrary to the geometry of the fringing, its *amplitude* is not stable in consecutive SCIENCE exposures since it strongly depends on the night sky conditions (as brightness, cloud coverage), the position on the sky and the airmass. If a reasonably good SUPERFLAT can be constructed from a sufficient number of SCIENCE frames obtained under stable sky conditions, a possible way to correct for fringes is the following:

1. Besides the large-scale sky variations not corrected for by SKYFLATs, the SUPERFLAT contains the fringe pattern as an additive component. Hence, the fringe pattern can be isolated from the SUPERFLAT by *FRINGE=SUPERFLAT−ILLUMINATION*.
2. Individual SCIENCE frames are corrected for fringing by *fringeSCIENCE=illumSCIENCE−f × FRINGE*. We assume that the fringe amplitude directly scales with sky brightness and *f* is calculated by *f=median(illumSCIENCE)/median(ILLUMINATION)*.

In the case of WFI@2.2m, this method removes very efficiently the low-level fringes in the R-band. Fringing is usually no longer visible and we estimate possible residuals well below 0.1%. In the case of I- and Z-bands, fringes can be suppressed to a level of about 0.1% of the sky amplitude if a very good SUPERFLAT can be constructed (see Fig. 6). If this is not the case, our approach may perform very poorly in reducing the fringe amplitude and we cannot propose a pipeline solution to the problem at this stage.

4.9. Guidelines for constructing calibration images

The success of the image pre-reduction heavily depends on the quality of calibration data at hand. In the creation of flatSCIENCE, we propagate the noise in masBIAS and masFLAT to our SCIENCE frames. As the noise in these calibration frames is of statistical nature we can diminish it by using as many images in the stacking process as possible. For the successful creation of a SUPERFLAT, and hence for the later quality of large-scale illumination and fringing correction, not only the number of images is important but it is essential that each pixel in the mosaic sees blank sky several times during SCIENCE observations. This suggests that the best observing strategy to achieve this goal is a large dither pattern between consecutive exposures (ideally it should be wider than the largest object in the field). In the case that the target occupies a significant fraction of the mosaic (as big galaxies or globular clusters for instance) the best strategy is to observe a neighbouring blank field for a short while if a SUPERFLAT and/or fringe correction is important. In the following we give some additional guidelines leading to good results in most of our own reductions:

- For the stacking process for masBIAS (masDARK) and masFLAT, about 15-20 individual images of each type should be acquired. As the noise in the final images is inversely proportional to the square root of the number of input images this is a good compromise between the number of exposures that can be taken during a typical observing run (in the case of SKYFLATs) and the noise reduction.
- The observer should find out the minimum exposure time required for the FLATs in order to suppress shutter effects. This can be done by taking DOMEFLATs or SKYFLATs with varying exposure time. One then normalises the flat fields and divides them. Any systematic residuals are then due to shutter effects, and one can determine the minimum usable exposure time. This is illustrated in Fig. 7 for the prime focus imager of the 4.2m William Herschel Telescope (WHT) at La Palma, which has a comparatively slow iris like shutter. Thus, images with short exposure times receive more light in the centre than in the corners. In any case, a SUPERFLAT of the data would recover any such shutter-related inhomogeneities in the FLAT, but we recommend to get everything correct from the beginning if possible. For newer or future WFIs such as OmegaCAM, shutters are constructed in which two blades move with constant speed from one side to the other exposing and covering each pixel of the detector array for exactly the same time (see e.g. Reif et al. 2004).
- As mentioned above, the creation of SUPERFLATs is not trivial in general. This is primarily an issue of which SCIENCE images to include. In the case of blue passbands, where only the smoothed ILLUMINATION image is used to correct for large-scale illumination gradients, it is safe to construct it from all SCIENCE frames of several, consecutive nights. Blue passbands are typically observed during periods of dark time, and hence the sky conditions are

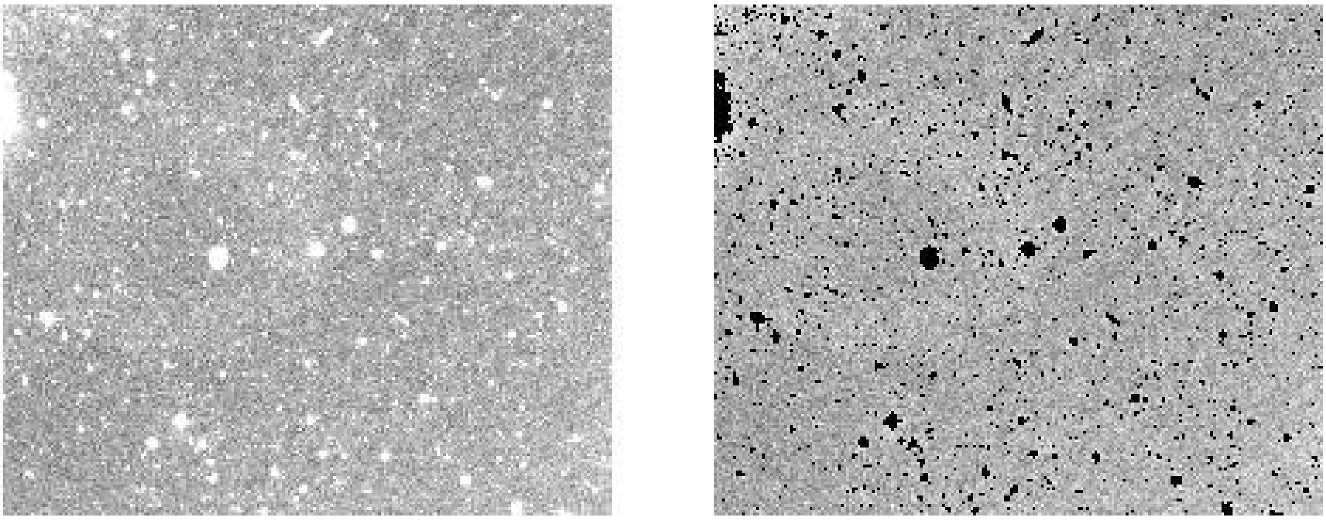

Fig. 4. For the creation of SUPERFLATs we mask pixels belonging to astronomical sources in the SCIENCE frames. The left panel shows a SCIENCE image before and the right panel after objects have been detected by `SExtractor` and flagged. We mask all structures having 50 contiguous pixels with 1.0 sigma above the sky-background. This helps to pick up also pixels of extended halos around bright stars. For images containing strong fringing we change the parameters to 7 pixels above 5 sigma as otherwise fringes would be detected and masked.

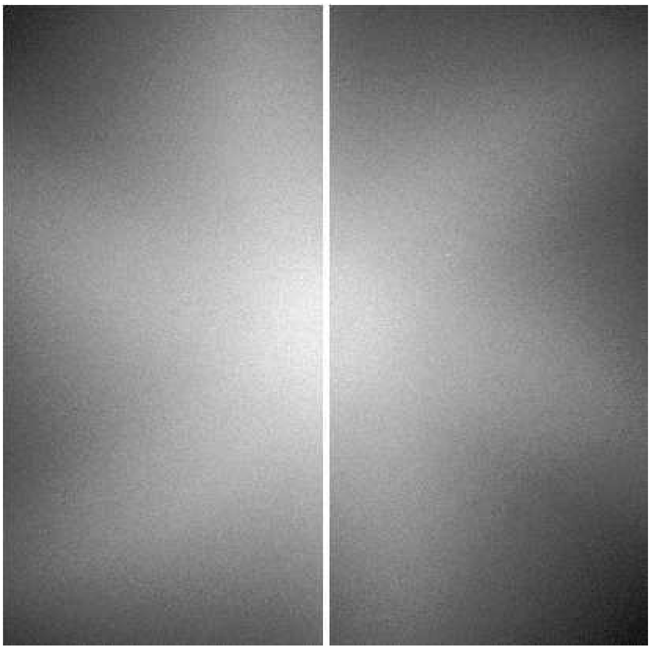

Fig. 7. Ratio of two normalised sky flats taken with the prime focus imager of the WHT (2 CCDs). The images were exposed for 0.3 and 1.0 seconds, respectively. One can see the way the iris shutter exposed and covered the detector again. The illumination difference between the brightest and faintest part in this representation is 20%. For the WHT prime focus imager it is recommended to expose for at least 2.0 seconds.

sufficiently stable. Also, with the moon down we observed no significant dependence of large-scale sky variations on telescope pointing. Hence, a robust SUPERFLAT can be obtained with a sufficient number of images during an observing run. This is no longer the case if we need to correct for strong fringing. We observe that our assumption of a sole dependence of the fringe amplitude on sky brightness typically fails when constructing the SUPERFLAT from SCIENCE frames of different sky positions. As the appearance of the overall pattern is very stable, this means either that the fringe amplitude is no longer a linear function of sky brightness or that the sky-background varies on scales significantly smaller than a CCD so that the scaling factor becomes dependent on CCD position. In addition, a similar behaviour is observed as a function of time, depending on the excitation of the OH^- night-sky emission lines. For our blank-field observations we obtain good results in the fringing correction if we observe the same target between 10 and 15 times with a large dither pattern within an hour (at WFI@2.2m with an overhead of about 2 minutes per image this can be achieved with 300 s exposures).

- If a very good SUPERFLAT could be constructed, flat-fielding results from our proposed method and from a pure SUPERFLAT application (i. e. using flat-fielding method 3 in Sect. 4.4) should be compared if flatness in SCIENCE images is crucial. Especially if only a small number of individual frames went into the construction of masFLAT, the direct SUPERFLAT approach often gives better results.
- Offsetting the telescope between the exposures is fundamental for high-quality, high-S/N mosaics. The *dither box*, i.e. the box that encloses all dither offsets, should be clearly larger than the gaps between the CCDs and the objects in the field. This has several advantages as compared to the *staring* mode (no offsets at all) or to the application of only small offsets.
  First, the CCDs in a multi-chip camera are fully independent from each other. They see the sky through

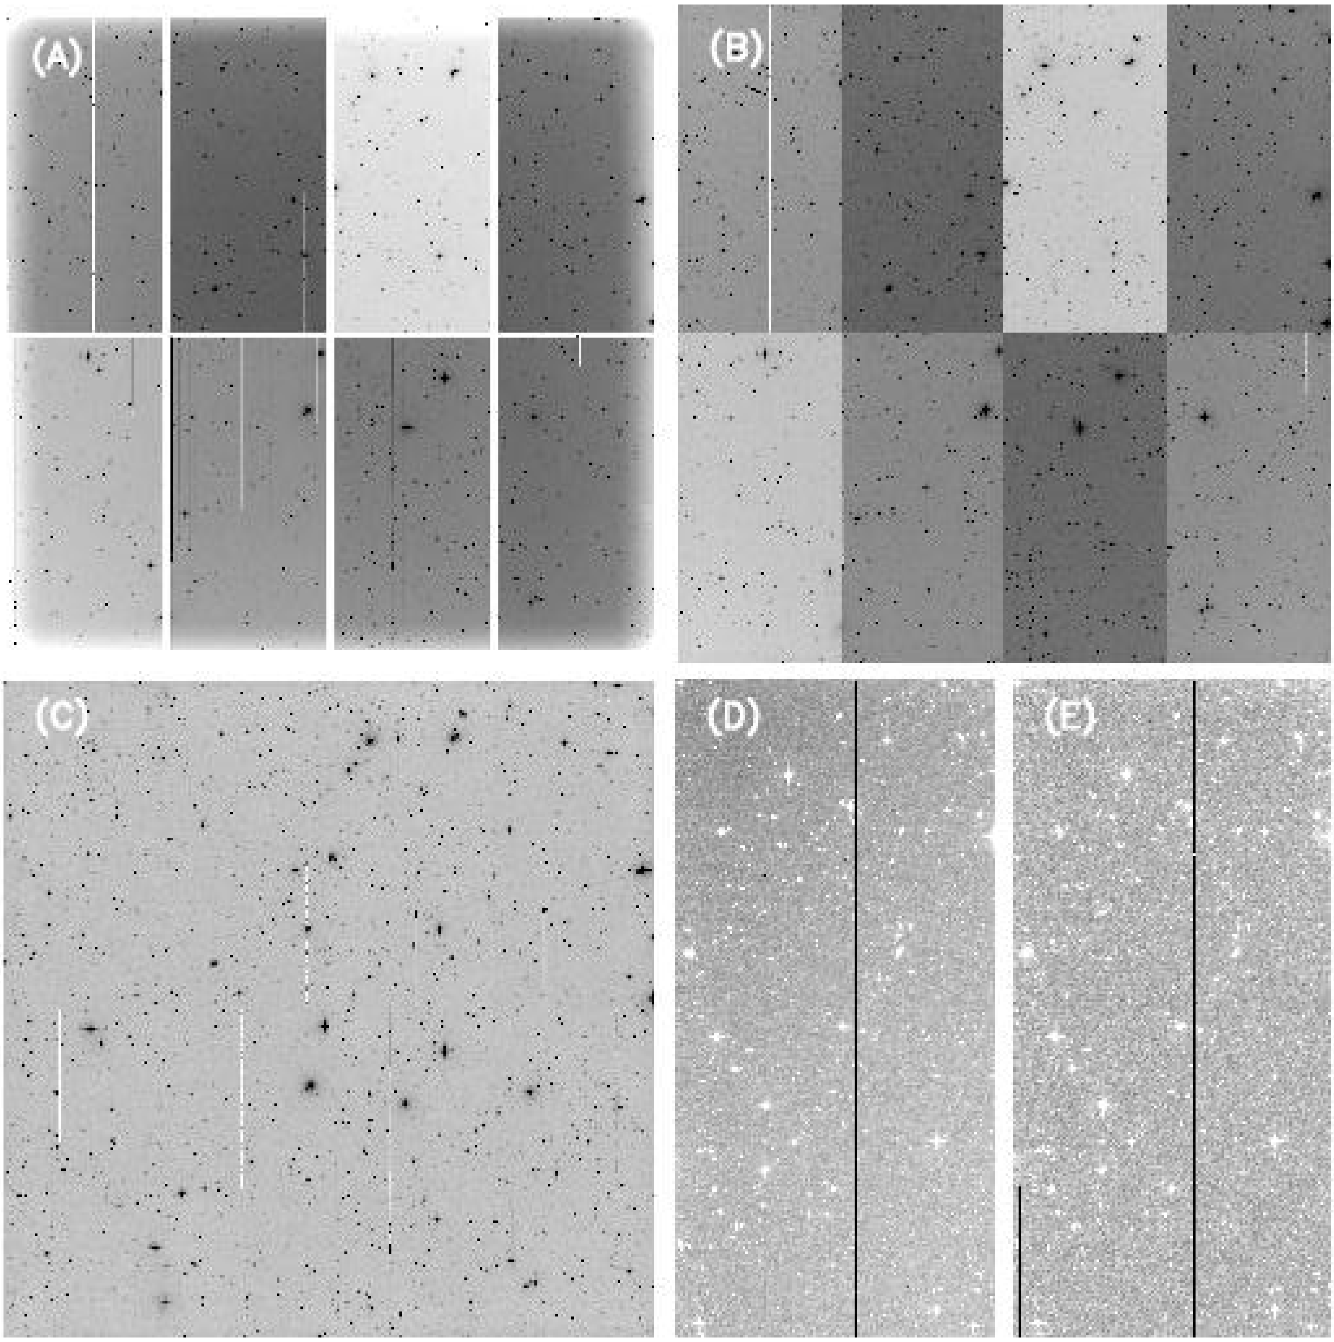

Fig. 5. The pre-reduction steps on a V-band exposure from WFI@2.2m: Panel (A) shows the raw image, panel (B) the result from applying masBIAS, trimming the CCDs and flat-fielding with the masFLAT. The high S/N masFLAT takes out small-scale variations but leaves large-scale residuals of up to 3% on the scale of the CCD [panel (D)]. These and the differences in the sky count-rate are removed after the application of ILLUMINATION giving a flatness of 1% over the entire mosaic in most of the cases as seen in panels (C) and (E).

different sections of the filter, and they have their individual flat-fields, gains and read-out noise. Choosing a wide dither pattern is the easiest way to establish an accurate global photometric and astrometric solution for the entire mosaic, based on enough overlap objects.

Second, the wide dither pattern allows for a significantly better superflattening of the data, since the objects do not fall on top of themselves in the stacks and thus for every pixel a good estimate of the background can be obtained. Besides, remaining very low-amplitude patterns in the sky background caused by improper flat-fields etc. do not add in the mosaic, but are averaged out. Thus, a wide dither pattern will lead to an improved sky background from which the S/N will benefit; see also Fig. 8.

Third, the object S/N in a mosaic with a (wide) dither pattern is significantly superior to one obtained in staring mode. The reason for this is that the master BIAS and the master FLAT are not noise-

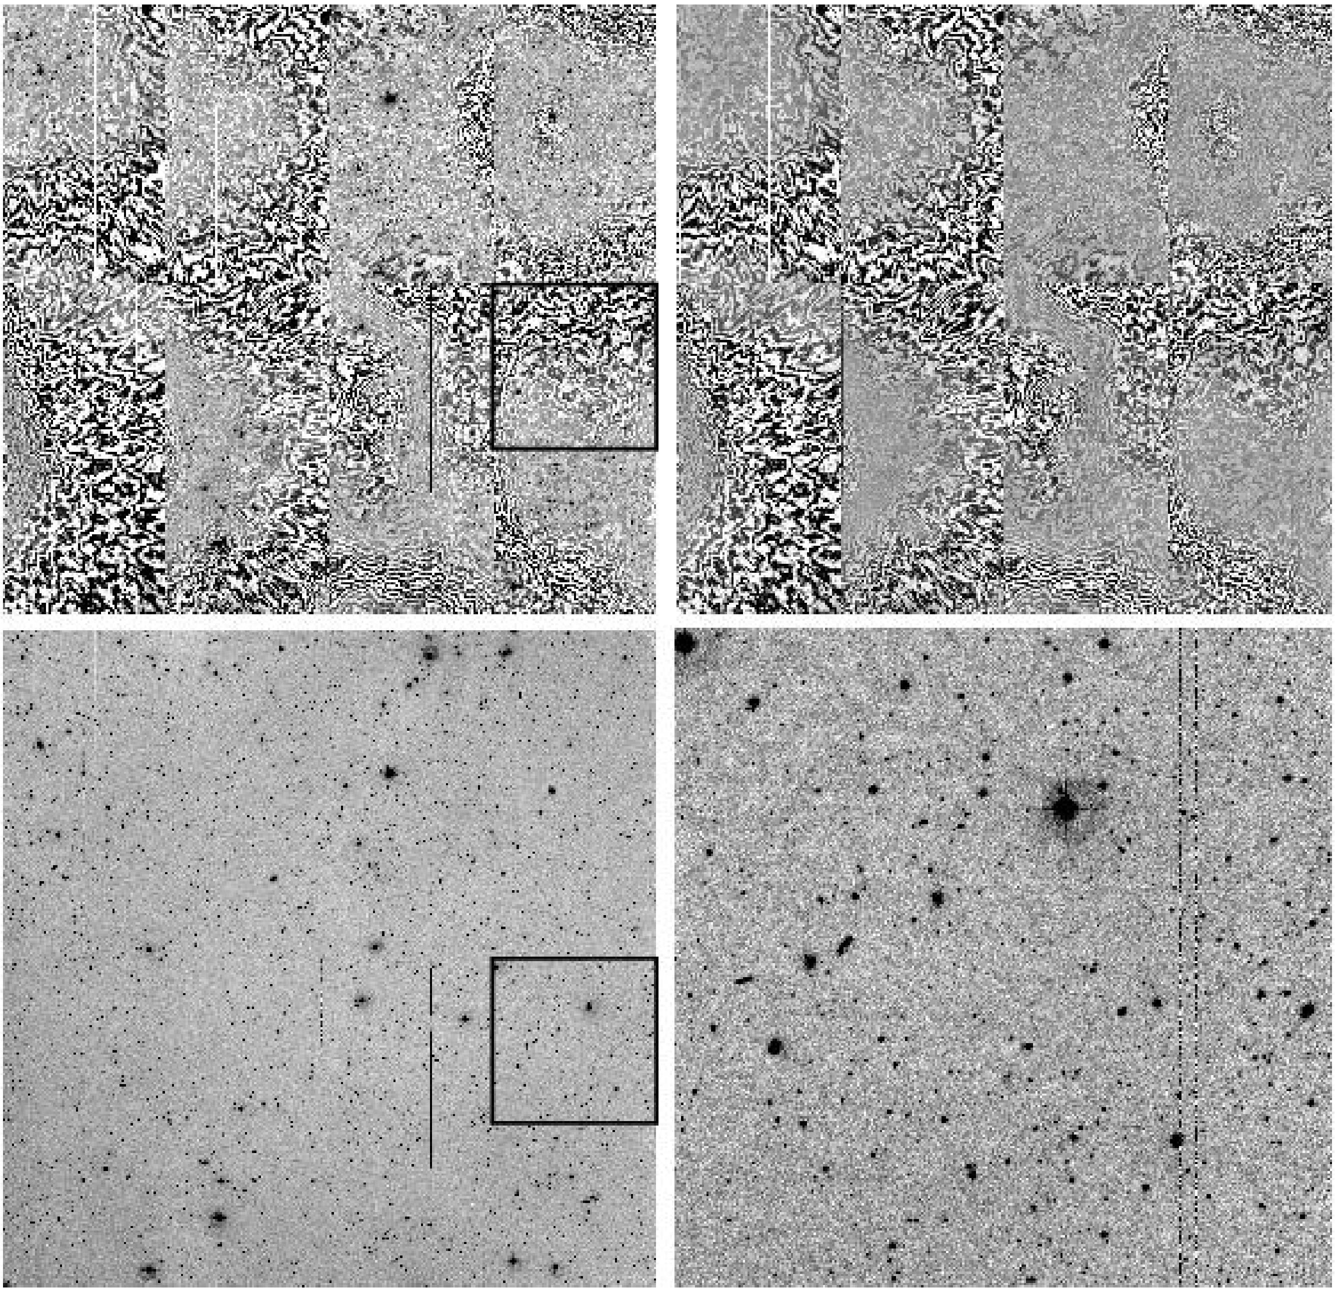

Fig. 6. Fringing correction in WFI@2.2m I-band data: For the shown correction, fifteen 300 s exposures have been obtained within 1.5 hours during stable, photometric conditions with a large dither pattern (3 arcmin in Ra and Dec) on a blank-field target. The upper left panel shows a SCIENCE frame before fringing correction, the upper right panel the extracted fringe pattern from the SUPERFLAT and the lower left panel the SCIENCE image after the correction. The lower right panel shows a zoom of a representative region (marked in the left two panels) after the correction. The residual fringes have an amplitude of about 0.1% of the night sky.

free, since they are created from a finite number of images. An identical copy of this noise is then created in each SCIENCE image during the pre-processing. If no dither offsets are applied, then the calibration noise in the N SCIENCE images is stacked on top of itself during the co-addition, and thus increases in the same way as the flux of the objects ($\propto N$) instead of averaging out ($\propto \sqrt{N}$). This is enforced by the use of an auto-guider, which keeps the sources in subsequent images exactly at the same pixel position. Depending on the ratio of the calibration noise and the noise in the uncalibrated images themselves, the effective exposure time of the mosaic can be very significantly reduced. This is especially true for observations gained under excellent (dark) conditions, since the sky noise in the SCIENCE images is then reduced, and the noise from the calibration frames becomes more dominant. One should therefore always aim at a sufficient number of calibration exposures,

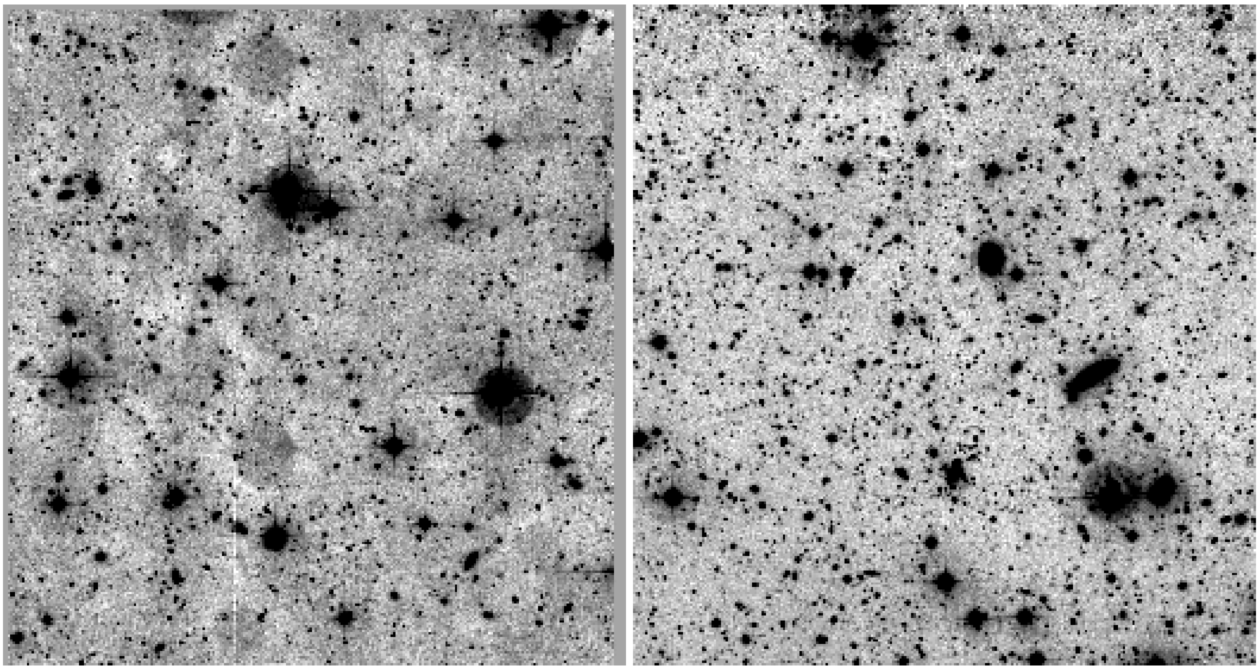

Fig. 8. Sky-background variations in a mosaic obtained from observations of several fields in staring mode (left), and for one obtained with a wide dither pattern (right). The contrast scaling for both mosaics is the same. Thus, science applications that require a very uniform background, e.g. the search for low surface brightness galaxies, profit in general substantially from large dither patterns.

and go for larger dither offsets. This holds for single-chip cameras, too.

4.10. Image Eyeballing and Manual Masking

After the pre-reduction we visually inspect all SCIENCE images to identify potential problems needing manual intervention:

- Faulty exposures (e.g. data taken out of focus) having passed the initial quality check (Sect. 4.3) are removed at this stage.
- Problems of the pre-reduction process are identified and necessary steps are iterated (this concerns mostly the selection of SCIENCE frames entering the SUPERFLAT).
- The most important aspect of this manual pass through all SCIENCE data is the identification of large-scale image defects that must not enter the final co-addition process and which have to be manually masked at this stage. This will be described in more detail in Sect. 6. See also Fig. 9.

4.11. Standard star processing and absolute photometric calibration

Besides the SCIENCE images we have to process STANDARD frames to obtain an absolute photometric calibration. A fully automated absolute photometric calibration up to the final co-added image is not yet implemented in our system. As a first step in this direction, the pipeline allows the photometric calibration of individual nights of a *Run*. In this section we describe how STANDARD exposures are pre-reduced and analysed to extract photometric parameters. If not stated otherwise, all magnitudes in this paper are given in the *Vega* system.

STANDARD exposures are reduced in a fashion similar to SCIENCE exposures. They are first overscan corrected, then the masBIAS is subtracted. Flat-fielding is done with the masFLAT generated in the pre-reduction of the SCIENCE exposures and the SUPERFLAT image created from the SCIENCE exposures is applied to the pre-reduced STANDARD exposures. Using the SUPERFLATs created from the SCIENCE images is important to ensure photometric homogeneity between SCIENCE and STANDARD exposures.

If necessary the FRINGE maps created from the SCIENCE images are rescaled and subtracted from the superflatted STANDARD exposures. This step is particularly difficult and often leads to non-satisfactory results in the I-band. The reason is that STANDARD exposures are usually taken at different times during the night (sometimes even in twilight), at different sky positions and various airmasses. This often leads to fringing patterns with intensities very different from those present in the SCIENCE exposures.

To generate object catalogues that can be used for photometric calibration bad pixel masks for the STANDARD images are created. They mark pixels not suitable for photometric measurements, such as bad pixels or cosmic rays. The creation of these masks (FLAG images henceforth) is described in more detail in Sect. 6.

Using the FLAGs object catalogues are created with `SExtractor` from the reduced STANDARD exposures. To match the found objects with Landolt (1992) or Stet-

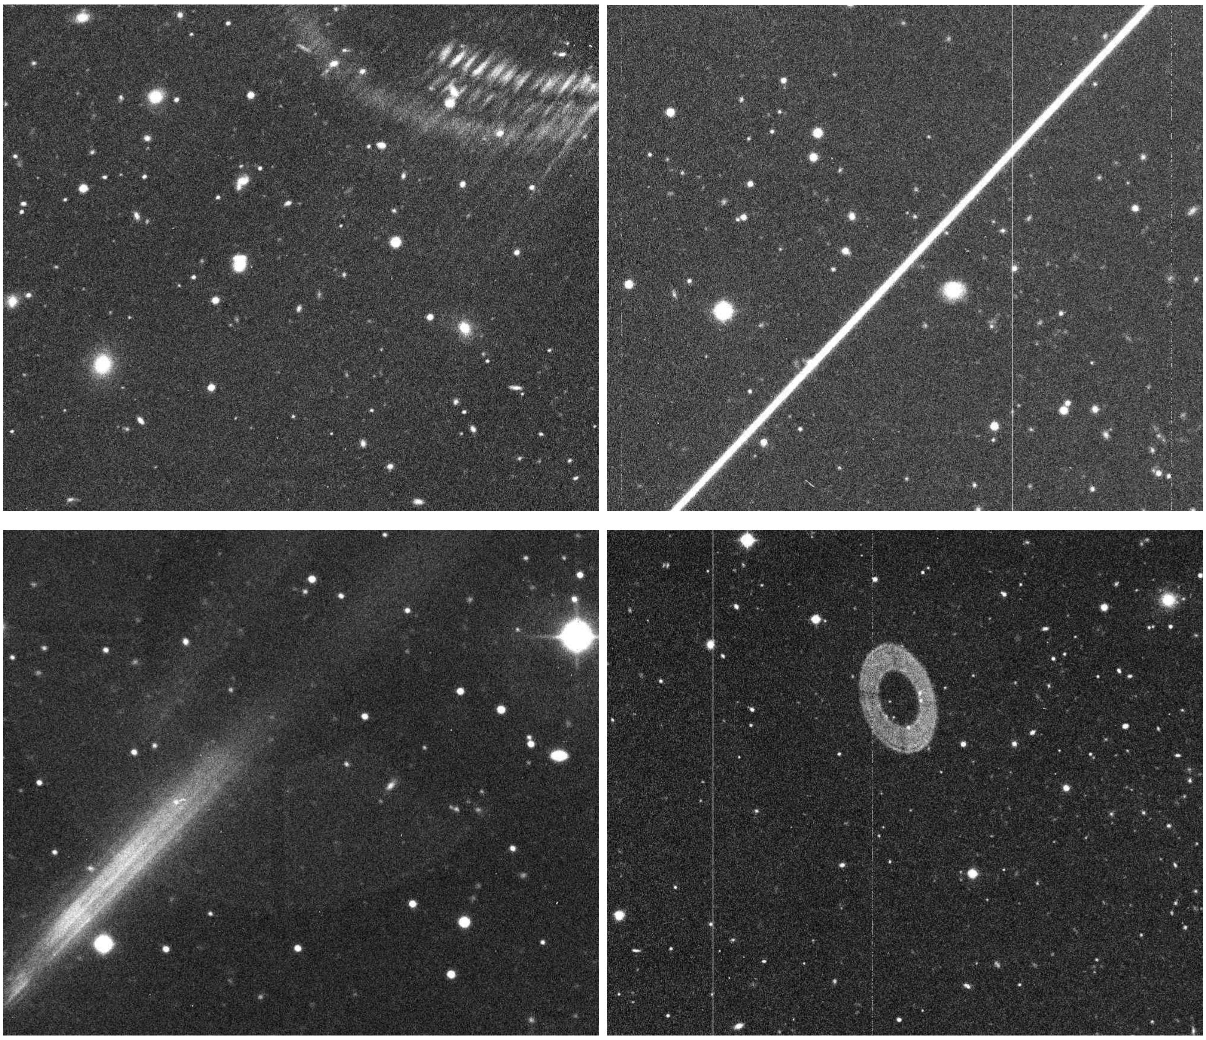

Fig. 9. Examples of large-scale defects: The upper right panel shows a satellite track, the other three panels various forms of reflections from bright stars inside and outside a WFI@2.2m field. All such features are typically unique to individual frames and do not appear in subsequent, sufficiently dithered exposures. Currently we have to register and to mask these defects by hand to prevent that they enter the final image co-addition (see Sect. 6). Masks are built using the *polygon* option of the `ds9` visualisation tool.

son (2000) standard star catalogues an astrometric solution for the STANDARD images has to be derived first. This is done with the LDAC astrom tools to match objects found in the STANDARD images to the USNO-A2 catalogue. The number of objects found in the STANDARD exposures is sufficiently high to derive a second-order astrometric solution separately for each image. This is enough to reliably merge standard star catalogues with our catalogues. All fluxes in our STANDARD fields are normalised to an exposure time of 1 s.

From the database of matched standard stars, the coefficients of the equation,

$$\mathrm{Mag} - \mathrm{Mag}_{\mathrm{inst}} = ZP + CE * \mathrm{airmass} + CT * \mathrm{CI}\ , \quad (1)$$

i.e. the zero point ZP, the extinction coefficient CE, and the colour term CT are determined, where Mag is the standard star's magnitude, $\mathrm{Mag}_{\mathrm{inst}}$ is the instrumental magnitude measured on the reduced standard frame [15], and CI is a colour index, e.g. $(B-V)$. This is done using a non-linear least-squares Marquardt-Levenberg algorithm with an iterative 3σ rejection, which allows rejected points to be used again if they are compatible with later solutions. As this algorithm is not guaranteed to converge, the iteration is aborted as soon as one of the following three criteria is true:

1. The iteration converged and no new points are rejected.
2. The maximum number of iterations (set to 20) is reached.

[15] we use the `SExtractor` MAG_AUTO parameter for this purpose. Although we do not know at which aperture we finally measure the standard stars magnitudes, MAG_AUTO turned out to be a reliable estimate for the total flux of bright sources. More importantly, we do not need to manually adapt the magnitude aperture for varying observing conditions, e.g. large seeing variations. We checked that our results with MAG_AUTO and appropriately chosen aperture magnitudes (`SExtractor` MAG_APER) are in excellent agreement.

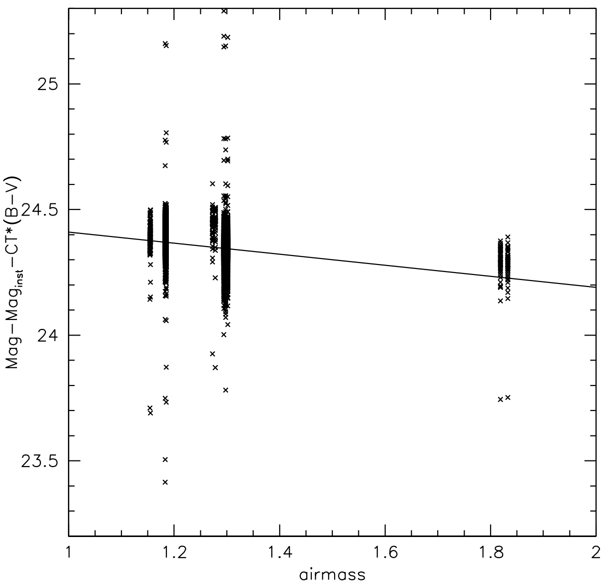

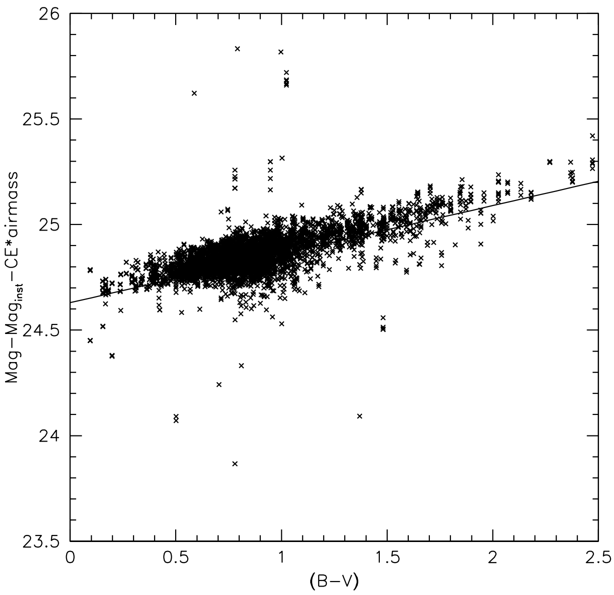

Fig. 10. The plots illustrate our zero point estimation in broad-band B during a clear night. Eleven standard field observations with a good airmass coverage were obtained giving a total of about 4000 individual standard star measurements (the Stetson standard star catalogue was used here). Points indicate the measurements, the lines the result of a three parameter fit of eq. (1).

3. More than a fixed percentage (set to 50%) of all points are rejected.

In a first step all three coefficients are fit simultaneously. However, in order to reliably estimate the extinction coefficient standard star observations must be spread over a range of airmasses. This is sometimes neglected by observers. To find an acceptable photometric solution in this case, the user can supply a default value for the extinction coefficient. The fit is then repeated with the extinction coefficient fixed at the user supplied value and the zero point and colour term as free parameters.

Although wide-field observations of Landolt/Stetson fields typically cover a wide range of stellar colours, the user can also supply a default value for the colour term. This is then used for a 1-parameter fit in which the zero point is the only free parameter.

In an interactive step the user can then choose between the 1-, 2-, and 3-parameter solution, or reject the night as non-photometric. The FITS headers of the frames belonging to the same night are then updated with the selected zero point and extinction coefficient or left at the default value −1.0, indicating that no photometric calibration for that frame is available.

We note that we perform the fit simultaneously for all mosaic chips. As discussed in Sect. 5.6 zero points of individual chips agree within 0.01-0.03 mag after sky-background equalisation. We do not take into account possible colour term variations between individual images that are expected due to slightly different CCD transmission curves. Notable differences can occur in the U and Z filters that are cut off by the CCD transmission.

Fig. 10 illustrates our photometric calibration which is implemented in `Python`.

4.12. *Runs* and *Sets*

As discussed in the beginning of Sect. 4 the pre-reduction is done on a *Run* basis usually containing observations from different patches of the sky. Before proceeding we need to split up the SCIENCE exposures of a *Run* into the *Sets* that have to be co-added later. By *Set* we mean the series of all exposures of the same target in a particular filter. This means that all the following reduction steps up to the final co-addition have to be done on each *Set* independently. We note that a *Set* can have been observed in multiple *Runs* so that all *Runs* containing images of a certain *Set* have to be processed at this stage.

5. Astrometric and photometric *Set* calibration

5.1. Astrometry

After the pre-processing, a global astrometric solution and a global relative photometric solution is calculated for all SCIENCE images. This is where the reduction of WFI data becomes much more complicated than the one for single-chip cameras.

In a first step, high S/N objects in each image are detected by `SExtractor`, and a catalogue of non-saturated stars is generated. Typically, we use all objects having at least 5 contiguous pixels with 5σ above the sky-background in the following analysis (these numbers may vary according to filter and exposure time; in the U-band for instance, we need to lower these thresholds in order to have enough sources to compare with a standard star catalogue). This usually gives us between 3

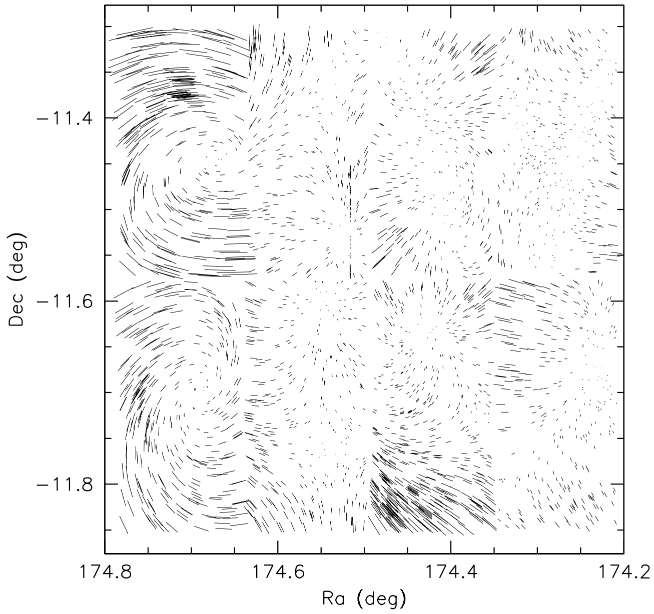

Fig. 11. Difference in object position between a single-shift approach and a full two-dimensional second-order astrometric solution for the WFI@2.2m. In other words, shown are the higher-order terms needed for matching the images to the sky. The patterns belonging to the left two CCDs are due to a rotation with respect to the mosaic. The maximum position difference in the plot is about six pixels, still a fairly small value compared to other telescope designs. It becomes clear that a single, global distortion polynomial for all 8 CCDs does not work. Instead, every image has to be treated individually.

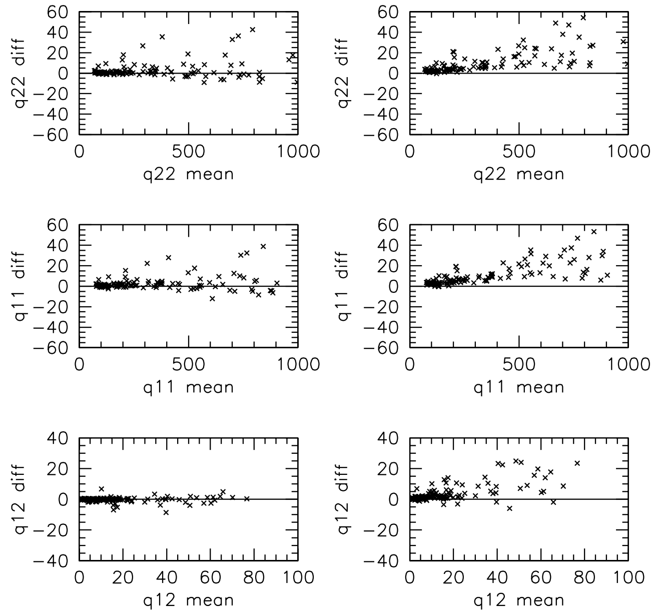

Fig. 12. Shown are measurements for $q_{ij}^{\rm diff}$ from 152 stars lying within one chip of 30 WFI@2.2m exposures. The quantities are plotted in arbitrary units. The co-addition was performed with `SWarp` using the LANCZOS3 resampling kernel, and the resampled images have been used for the calculation of $q_{ij}^{\rm mean}$ (see Sect. 7). All stars have at least 10 contiguous pixels with 10σ above the sky-background in the individual images and hence all the measurements involved have high S/N. The left panels show the result for $q_{ij}^{\rm diff}$, where the co-added mosaic was stacked according to the astrometric solution provided by `Astrometrix`. We checked that the outliers showing a higher value in the co-added image are uniformly distributed over the images. For the right panels we shifted all individual images with a random offset within an image pixel before co-adding them. This mimics a co-addition by a simple integer pixel shift and we see a clear, systematic offset for all three components.

and 6 objects per square arcmin in high-galactic latitude empty field observations. Based on a comparison with the USNO-A2 (see Monet et al. 1998) astrometric reference catalogue (or any other reference catalogue), a zero-order, single shift astrometric solution is calculated for each image. For a single-chip camera with a small field-of-view such an approach is often sufficient, but it no longer holds for multi-chip cameras with their large field-of-view. In this case, the CCDs can be rotated with respect to each other, tilted against the focal plane, and in general cover areas at a distance from the optical axis where field distortions become large. Fig. 11 shows the difference between a zero order (single shift with respect to a reference catalogue) and a full astrometric second-order astrometric solution per image. From this figure it is obvious that the simple shift-and-add approach will not work for the entire mosaic. The issue is further complicated by the gaps between the CCDs and large dither patterns that are used to cover them. Thus, images with very different distortions overlap. In addition, due to the large field-of-view, one must take into account that the spherically curved sky is mapped into the flat tangential detector plane.

In the second step, Mario Radovich's `Astrometrix`[16] package is used to determine third-order polynomials for the astrometric solution of each image. This corrects for the aforementioned effects, and thus allows for the proper mosaicing in the later co-addition process. For this purpose all high S/N objects (stars and galaxies) detected in the first step are matched, including those from the overlap regions. The latter ones are most important in establishing a global astrometric (and photometric) solution, since the accuracy of available reference catalogues such as the USNO-A2 with an rms of about 0$''$.3 is insufficient for sub-pixel registration. Thus the astrometric solution is determined from the data itself instead of the reference catalogue, which is usually based on data taken with a much smaller angular resolution than the images that are processed. The reference catalogue is used only to fix the solution with respect to absolute sky coordinates within its stated rms. A wide dither pattern is required to compute a reliable global astrometric solution, so that different regions of the CCD are mapped by overlapping sources and the astrometric solution is properly constrained. For a more detailed de-

[16] http://www.na.astro.it/∼radovich/WIFIX/astrom.ps

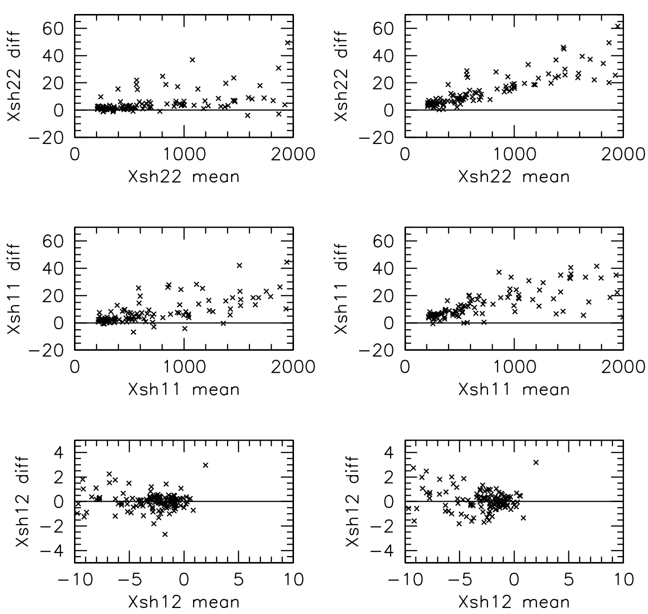

Fig. 13. The same as Fig. 12 for $X^{\rm sh}$ in arbitrary units. The agreement between the measurement in the mosaic and the expectation is not as good as for q and $X^{\rm sm}$ (see Fig. 14) and small systematic offsets up to 1% are seen in the 11 and 22 components.

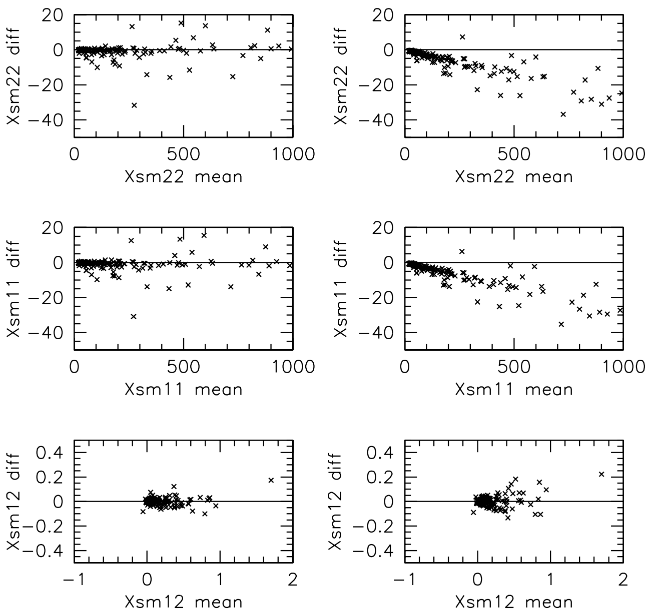

Fig. 14. The same as Fig. 12 for $X^{\rm sm}$ in arbitrary units. For this quantity, the agreement between the values in the mosaic and the predictions is excellent for all three components.

scription of the `Astrometrix` tool see Radovich (2002) and McCracken et al. (2003).

5.2. Accuracy of our Astrometric calibration

In the following discussion we mainly focus on the *internal* astrometric accuracy obtained in R-band exposures with WFI@2.2m. The astrometric calibration ensures that we can produce a deep mosaic out of individual images. The accuracy with which the images have to be aligned strongly depends on the scientific application in mind. In most cases, the final goal is the measurement of moments M^{mn} from the light distribution $I(\theta)$ of objects, that is:

$$M^{mn} = \int \mathrm{d}^2\theta\, \theta_1^m \theta_2^n W(\theta) I(\boldsymbol{\theta}); m, n \in \mathcal{N}_0$$

where $W(\theta)$ is some weighting function, $\theta = |\boldsymbol{\theta}|$ and $(\theta_1, \theta_2) = (0,0)$ is the centre of the object, which is chosen to be the point where the weighted dipole moments are zero. In our weak lensing studies, we have to estimate accurately moments up to the fourth order. The shape of an object is constructed out of the second moments q_{ij}, i.e.;

$$q_{ij} = \int \mathrm{d}^2\theta\, W(\theta)\, \theta_i \theta_j I(\boldsymbol{\theta}).$$

The following two fourth-order moments are important for correcting object shapes for PSF effects in weak lensing studies (see Kaiser et al. 1995):

$$X^{\rm sh} = \int \mathrm{d}^2\theta \begin{bmatrix} 2W\theta^2 + 2W'\eta_1^2 & 2W'\eta_1\eta_2 \\ 2W'\eta_1\eta_2 & 2W\theta^2 + 2W'\eta_2^2 \end{bmatrix} I(\boldsymbol{\theta})$$

and

$$X^{\rm sm} = \int \mathrm{d}^2\theta \begin{bmatrix} \eta_3 + 2W''\eta_1^2 & W''\eta_1\eta_2 \\ W''\eta_1\eta_2 & \eta_3 + W''\eta_2^2 \end{bmatrix} I(\boldsymbol{\theta}),$$

where $\eta_1 := (\theta_1^2 - \theta_2^2)$, $\eta_2 := 2\theta_1\theta_2$, $\eta_3 := W + 2W'\theta^2$ and W' denotes the derivative with respect to θ^2. As the quantities under consideration are linear in the light distribution $I(\boldsymbol{\theta})$, the measurement in the mosaic can be predicted from the measurements in the individual images (if a linear image co-addition is performed). If the mosaic is constructed by a straight mean of N input images, the value $q_{ij}^{\rm mosaic}$, measured in the co-added image has to be equal to:

$$q_{ij}^{\rm mean} = \frac{1}{N}\sum_{k=1}^{N} q_{ij}^k,$$

where the q_{ij}^k are the moments in the individual images. In the following we choose $W(\theta) = 1/(2\pi r_{\rm g}^2)\exp(-|\theta|^2/(2r_{\rm g}^2))$ with $r_{\rm g} = 3$ image pixels in all calculations. A systematic deviation of the quantity $q_{ij}^{\rm diff} := q_{ij}^{\rm mosaic} - q_{ij}^{\rm mean}$ (and in corresponding expressions for $X^{\rm sh}$ and $X^{\rm sm}$) from zero is hence a good indicator that the astrometric solution should be improved. Results of this test for a deep R-band mosaic, constructed out of 30 exposures with sub-arcsecond seeing, are shown in Figs. 12, 13 and 14. The agreement between the predictions and the measurements is in general very good, only the results of $X_{\alpha\beta}^{\rm sh}$ may indicate a small systematic offset in the co-added image.

For the positional accuracy of overlap sources (i.e. the alignment accuracy of first-order moments of the light distribution) `Astrometrix` formally estimates an rms accuracy of $1/10 - 1/5$ of a WFI pixel ($0\farcs02 - 0\farcs04$) over the whole field-of-view and we give an example in Fig. 15. This accuracy is consistent with results obtained by Alcalá et al. (2002) for WFI@2.2m observations.

However, the quoted accuracies strongly depend on filter, instrument and the observing strategy, i.e. the number of bright stars that can be extracted, the kind of distortions and the dither pattern. For instance, Radovich et al. (2004) obtain an internal alignment accuracy of about 1/3 of a pixel ($0\farcs07$) for WFI@2.2m U-band observations and McCracken et al. (2003) 1/4 of a pixel ($0\farcs05$) for B, V, R, I observations with the CFH12K@CFHT instrument. In Fig. 16 we compare astrometric solutions from different bands in one of our own WFI@2.2m survey fields and note that for this instrument we are able to align different broad-band filters to sub-pixel accuracy with independent astrometric calibrations in the different bands.

At the end of this section we compare in Fig. 17 absolute astrometric positions in our co-added images, which are tied to the USNO-A2 catalog, with objects in the more accurate, but sparser UCAC2 standard star catalogue (see Zacharias et al. 2004). From this analysis we conclude that most of our fields have an absolute astrometric inaccuracy on the order of $0\farcs5$ in Ra and Dec.

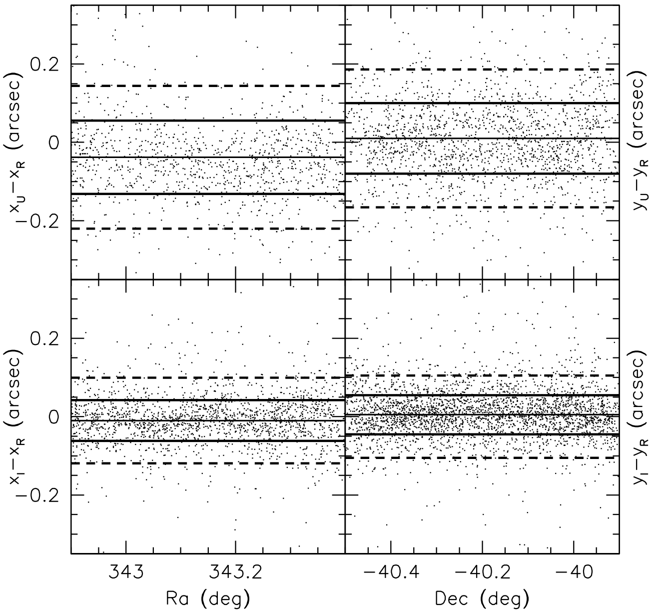

Fig. 16. Residuals of object positions in the co-added images of different broad-band filters (I, R and U) for one of our WFI@2.2m survey fields are shown. We note that all astrometric calibrations in the different filters were done *independently* from each other. Each band contains at least 30 exposures obtained with a large dither box of about $3\farcm0$. Compared to R and I, individual U band observations contain significantly fewer high S/N sources suitable for astrometric calibration and hence we expect a less accurate alignment for this filter. The plot shows in the lower panels that the astrometric systems for R and I are in very good agreement. The absolute positional offset between these two bands is $\Delta\mathrm{Ra} = -0\farcs01; \Delta\mathrm{Dec} = 0\farcs005$. Solid (dashed) lines indicate areas containing 68% (90%) of all points. They are at $\Delta\mathrm{Ra} = \pm0\farcs052(\pm0\farcs11); \Delta\mathrm{Dec} = \pm0\farcs051(\pm0\farcs10)$. The U band has a larger offset ($\Delta\mathrm{Ra} = -0\farcs038; \Delta\mathrm{Dec} = 0\farcs011$) and an rms scatter of about half a pixel with respect to R (upper panels). The lines are located at $\Delta\mathrm{Ra} = \pm0\farcs094(\pm0\farcs18); \Delta\mathrm{Dec} = \pm0\farcs091(\pm0\farcs18)$ with respect to the offset centre. We note that all bands are aligned to sub-pixel accuracy. For clarity the plots show only one point out of five.

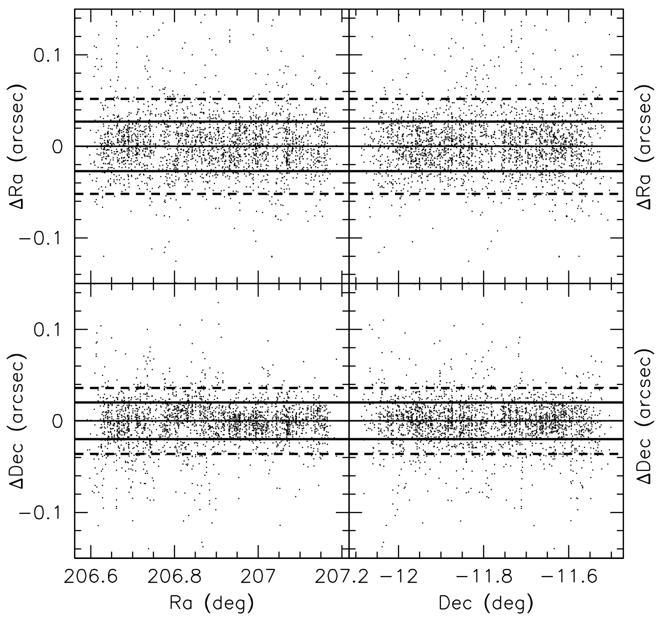

Fig. 15. The positions of sources in 28 individual R-band WFI@2.2m exposures after astrometric calibration are compared with the locations of their counter-parts in the co-added mosaic. The observations were obtained with a large dither box of about $3\farcm0$. The thick solid lines mark the region containing 68% of all points and are at $\Delta\mathrm{Ra} = \pm0\farcs027; \Delta\mathrm{Dec} = \pm0\farcs02$. Dashed lines show the corresponding area for 90% of all points ($\Delta\mathrm{Ra} = \pm0\farcs052; \Delta\mathrm{Dec} = \pm0\farcs036$). Only one out of five points is shown for clarity of the plot.

5.3. Photometry

Besides the astrometric alignment, special care has to be taken in the photometric calibration. First, we usually want to co-add data obtained under different photometric conditions (for instance if a target is observed during different nights of an observing run). Second, the individual images can have intrinsically different zero points[17]. It is possible to arrive at an *internally consistent* photometric system for all exposures by comparing instrumental magnitudes of overlap sources. The following procedure of a relative photometric calibration is very similar

[17] The sky-background equalisation described in Sect. 4.6 should remove zero point differences within an exposure but we do not rely on this assumption here.

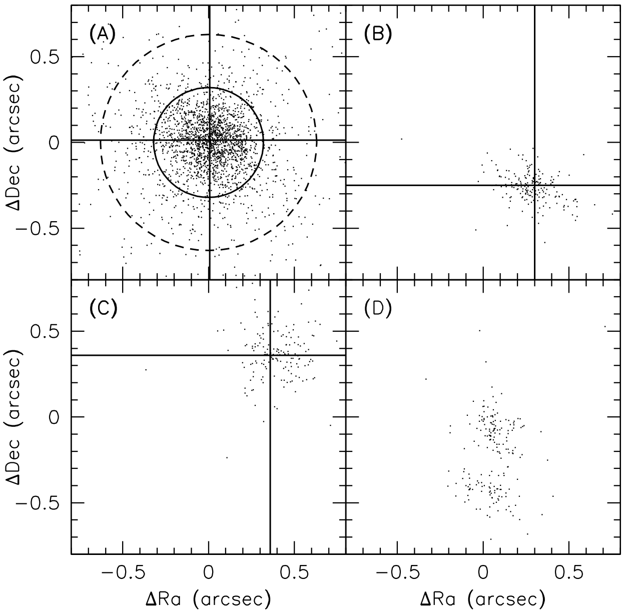

Fig. 17. Shown are positional residuals from sources in astrometrically calibrated and co-added WFI@2.2m mosaics compared to objects in astrometric standard star catalogues. Panel (A) displays residuals with respect to the USNO-A2 catalogue which is the reference system for our calibrations. The solid and dashed circles at 0.″32 and 0.″63 enclose 68% and 90% of all objects. The cross marks the centre of the distribution and lies at ΔRa = 0.″007 and ΔDec = 0.″013. The other three panels show residuals of various survey fields with respect to the UCAC2 catalogue. This catalogue contains about 10 times fewer objects than USNO-A2 but its absolute positional errors are 4 to 10 times smaller (20-70mas) which allows a good estimate for the absolute astrometric accuracy of our co-added images. For many fields we observe a shift of the two systems of up to 0.″5 in Ra and Dec as seen in panel (B) with ΔRa = 0.″29 and ΔDec = −0.″25 and panel (C) with ΔRa = 0.″36 and ΔDec = 0.″36, but also more complex residual patterns occur as seen in panel (D). In this case our survey field probably crosses the border of two photographic plates from USNO-A2.

to that described in Koranyi et al. (1998). It is implemented in the LDAC `relphotom` programme. Having a full astrometric solution and hence an accurate table of paired objects of all SCIENCE images at hand, a relative photometric solution is straightforward. Given two overlapping images k and j, consider all $i = 1...N$ common objects and calculate the mean deviation of magnitudes K and J

$$M_{k,j} := \frac{\Sigma_i W_i (K_i - J_i)}{\Sigma_i W_i} \,, \tag{2}$$

with $W_i = (\sigma_K^2 + \sigma_J^2)^{-1}$, where the σ are the measurement errors of the corresponding magnitude estimates. Objects deviating in $K_i - J_i$ more than a user-defined threshold are excluded from the following statistics. The relative zero points ZP_l for all N overlapping images are estimated by a χ^2 minimisation with respect to ZP_k,

$$\chi^2 = \sum_{k,j}^{N} \left[M_{k,j} - (ZP_k - ZP_j)\right]^2 \,. \tag{3}$$

In addition we demand that $\sum_k ZP_k = 0$ to make the system non-degenerate. We test the robustness of our photometric algorithms with simulations and on real data in the next sections. We note however that our relative photometric calibration depends on two important assumptions that may not be fulfilled:

- We assume that the relative zero points are constant on the scale of each image. As reported by Koch et al. (2004), zero point variations of up to ≈ 0.06 mag within a WFI@2.2m image can occur due to non-uniform illuminations (e.g. scattered light or reflections inside the instrument). Manfroid et al. (2001) show that these deviations are not corrected for by standard flat-field techniques (such as the application of SKYFLATs or SUPERFLATs). Indeed, they argue that the aim of a flat sky-background is incompatible with the aim of homogeneous photometry on the scale of an image. Similar effects have been observed with the CFH12K@CFHT instrument (see Magnier & Cuillandre 2004). In other words, the main reason for the observed variations of the photometric zero point is the mixing of multiplicative (flat-field) and additive effects (scattered light, reflections) in the flat-fields. The division by such a flat then leads to a homogeneous surface brightness, but introduces the observed variations in the photometric zero point. A solution would be to disentangle the multiplicative and additive component. At the moment it is not clear how this could be achieved in the general case.
- We neglect that objects of different spectral type can be affected differently by varying photometric conditions (e.g. absorption by clouds). At this stage we have not investigated this effect.

For the moment we do not take into account these effects in our photometric calibration.

For a full photometric calibration of our SCIENCE images we need to include the results of our standard star calibration (Sect. 4.11) in the estimation of our relative zero points. We are currently implementing and testing this.

5.4. Simulations for investigating the robustness of the photometric calibration

The simulations described in this section aim at quantifying the following issues of our photometric calibration:

- How accurately can we recover relative photometric zero points between individual images?
- We subtract the sky-background from our images before they are finally co-added (see Sect. 7 for details on the procedure). Does this procedure introduce any systematic bias in the flux measurements of faint objects?

To this end we simulated WFI exposures with Emmanuel Bertin's `skymaker`[18] (see Erben et al. 2001) and processed them through our pipeline in exactly the same way as real data. In this way we include all possible uncertainties from measurements, possible pixelisation effects etc. in our simulations. The simulations have the following characteristics:

- We created 25 exposures of 8 chips in the layout of the WFI@2.2m telescope (i.e. a pixel size of 0.″238 and a chip size of 2k × 4k). Each exposure mimics a 500 s integration in an R broad-band filter. The images do not contain any flat-field or fringe residuals but resemble perfectly pre-reduced data.
- Each exposure is put randomly within a $2' \times 2'$ dither box. No geometric distortions have been introduced.
- Each exposure has a random image seeing between 0.″9 and 1.″1 (no PSF anisotropy) and a random sky R-band background between 20.0 and 20.5 mag/arcsec^2.
- Each image is assigned a random magnitude zero point between 24.9 ± 0.3, i.e. there is no correlation of zero point with exposure, extinction or gain.

The images were processed through our pipeline starting at the astrometric calibration up to the final mosaic. The coaddition was done with `SWarp` using the LANCZOS3 kernel (see Sect. 7). In addition we simulated with `skymaker` a single 10k × 10k exposure with a formal integration time of 25 × 500 s (skymaker mosaic in the following) for flux comparisons with the co-added mosaic. The accuracy of the zero point estimation is shown in Fig. 18. We see that our procedure recovers the input values with a formal uncertainty of 0.002 mag. In Fig. 19 we show the comparison of flux measurements from the co-added mosaic with the skymaker mosaic. No systematic deviations in the magnitudes up to the faintest objects are visible. We conclude that our procedures accurately recover relative zero point deviations between images and that our mosaics do not suffer from systematics in estimating magnitudes. Errors connected to the algorithms are, at least for WFI@2.2m, small compared to systematic errors (variations of about 0.05 mag within an image) due to possible CCD illumination problems (see above).

5.5. Photometric measurements in NGC 300

To finally quantify the quality of our current photometric calibration, we compare flux measurements in the field of NGC 300 (see Schirmer et al. 2003) with previously published secondary standards by Pietrzyński et al. (2002). PSF photometry was carried out on co-added B (exposure time: 39 600 sec) and V (exposure time: 37 440 sec) WFI@2.2m images of NGC 300 by using the DAOPHOT task (Stetson 1987) implemented in IRAF Version 2.12.2. We used about 100 out of 390 secondary photometric standard stars that were published by Pietrzyński et al. (2002) to uniformly cover our field-of-view. Areas, in which crowding became significant or where estimations of the sky-background became difficult (galaxy centre, spiral arms, saturated stars), were excluded. Fig. 20 shows the spatial distribution of the chosen standard stars in the co-added image. The galaxy-centre of NGC 300 coincides with the field centre. The photometric standards published in Pietrzyński et al. (2002) were observed in four fields with the Warsaw 1.3 m telescope at the Las Campanas Observatory.

[18] the TERAPIX tool for the simulation of astronomical images

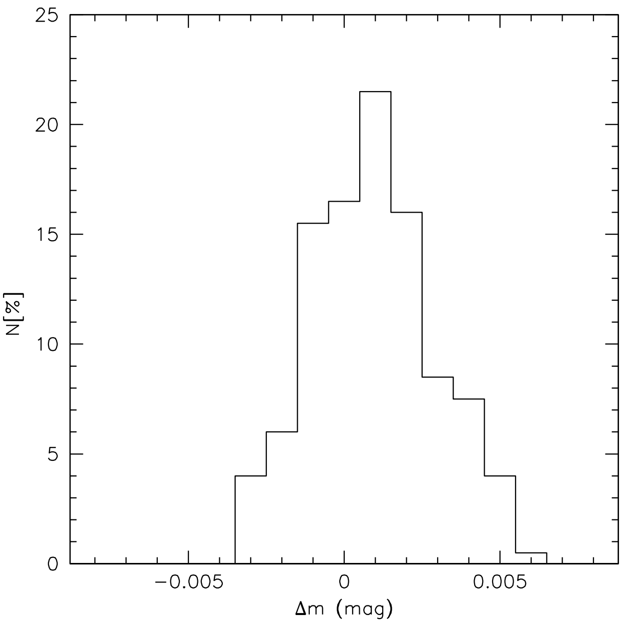

Fig. 18. The figure shows the error distribution of relative photometric zero point recovery in Monte Carlo simulations. The x-axis displays the difference Δm (in magnitudes) between input and recovered relative zero point, the y-axis the recovery probability. Our algorithm recovers the input values formally with $\Delta m = 0 \pm 0.002$. See the text for more details on the simulations.

Using the data we solved the following transformation where we neglect the airmass-term which is constant over the field-of-view:

$$b \;=\; B + b_1 \cdot (B - V) + b_2 \tag{4}$$

$$v \;=\; V + v_1 \cdot (B - V) + v_2. \tag{5}$$

Capital letters represent standard star magnitudes in B and V respectively whereas lower case letters indicate the measured instrumental magnitudes. Values of $b_{[1,2]}$ and $v_{[1,2]}$ have been determined by solving these equations for the standard stars using the Levenberg-Marquardt algorithm implemented in the IRAF routine INLFIT. This leads to values of $b_1 = -0.28 \pm 0.007$, $b_2 = 0.69 \pm 0.006$, $v_1 = 0.078 \pm 0.005$ and $v_2 = 0.89 \pm 0.005$. In Fig. 21 we plot the residuals of this fit as a function of B-band (left-hand side) and V-band (right-hand side) magnitude and see that there is (to a first approximation) no significant tendency for the residuals to correlate with

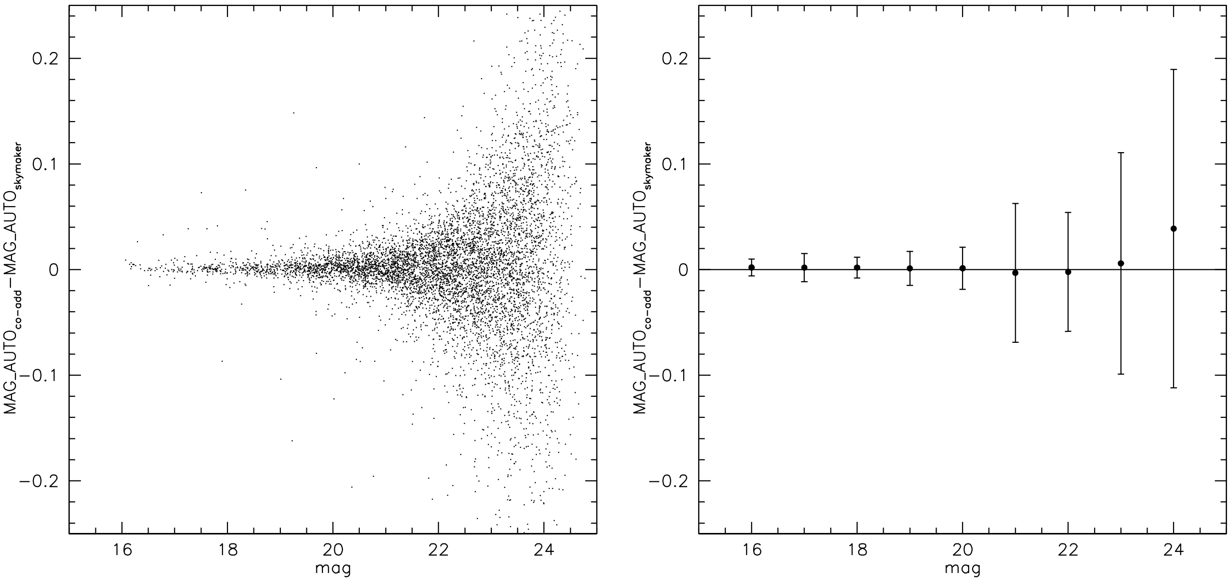

Fig. 19. Investigation of possible flux measurement biases in our mosaics: We simulated 25 dithered WFI exposures with `skymaker` and processed them through our pipeline in exactly the same manner as real data. In addition we created one large skymaker mosaic not undergoing any pipeline processing. The figure shows the difference of magnitude estimates in these two images, i.e. it mainly reveals possible biases due to inaccurate sky subtraction in the mosaic. The left panel shows the raw differences, the right panel bins with a width of 1 magnitude. We notice that our processing does not introduce systematic biases up to the faintest magnitudes. See the text for more details.

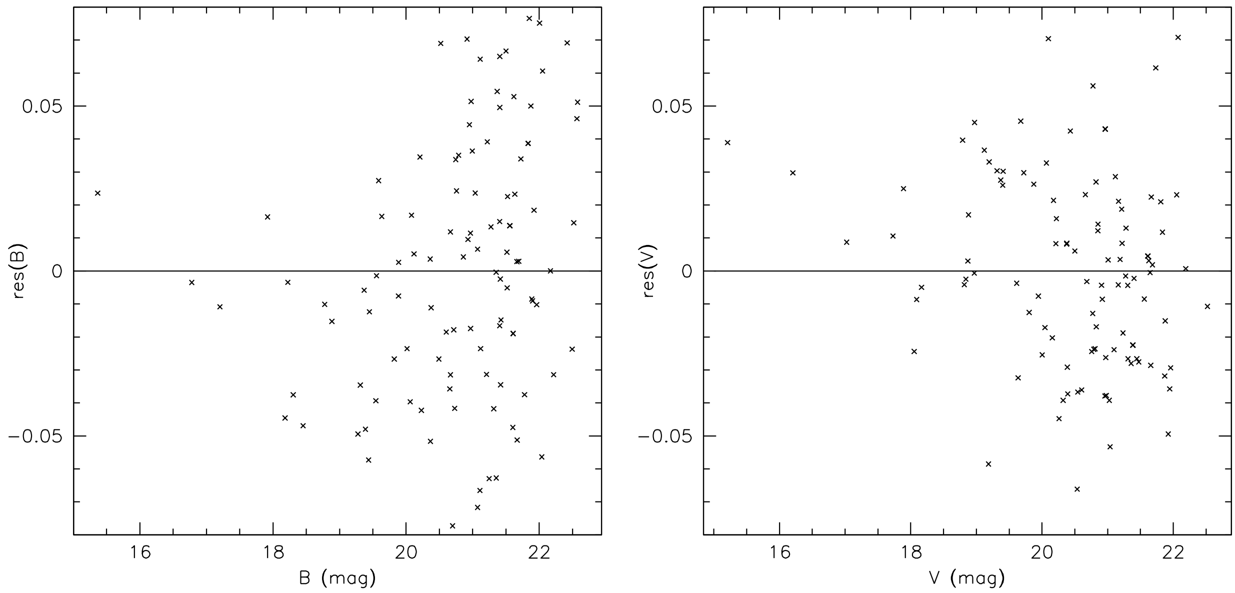

Fig. 21. Residuals vs. magnitude plots of 107 chosen standard stars. The left panel shows the results for the B-band, whereas the right panel shows those for the V-band.

magnitude[19]. Calculating the standard deviation σ for these fits we get $\sigma(B) = 0.037$ and $\sigma(V) = 0.029$ with maximum residuals of about 0.07 magnitudes.

This means the errors and/or uncertainties introduced by a global solution for the majority of stars are smaller than 0.04 magnitudes and are therefore negligible for most photometric studies taking into account that they are of the same order as the uncertainties for some of the secondary photometric standards.

We also tested the global solution in relation to the different chips. For this, we fitted the global solution separately to areas where mostly images of a particular chip overlap and checked the standard deviation of these fits. The results are presented in Table 3 and it becomes obvious that the standard deviation for a single chip dif-

[19] this especially holds for stars fainter than mag=19.

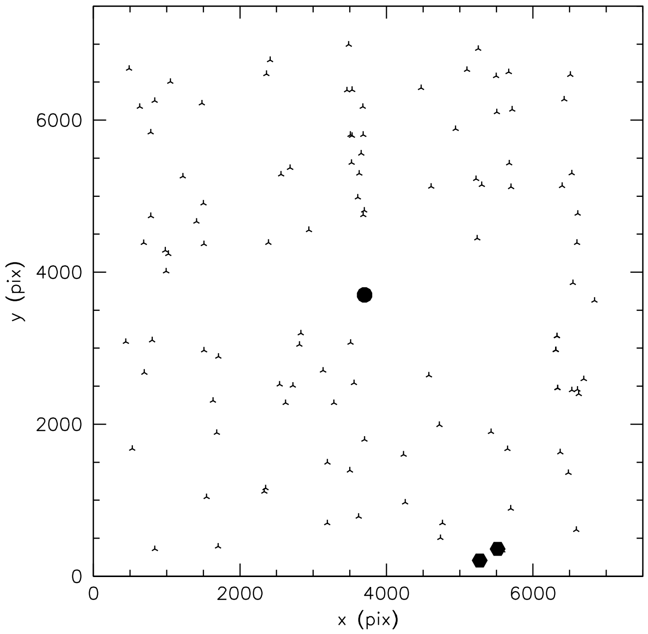

Fig. 20. Positions of the chosen standard stars in the co-added V-band image are shown. Crosses denote the positions of the secondary standards from Pietrzyński et al. (2002), the big dot the centre of NGC 300 and the filled hexagons two photoelectric standards from Graham (1981).

fers from the standard deviation of the global solution only by a constant smaller than 0.005. Due to this, our global solution can be regarded as sufficient and therefore shows that our photometric calibration is capable of removing chip-to-chip variations in a more than satisfying manner. This has also to be seen in the context that the observations of the very extended NGC 300 galaxy were spread over 34 nights in six months, encompassing very different sky conditions.

Finally we have performed a check of our transformation with two photoelectric standards published by Graham (1981) which are not saturated in our images. Applying the global solution to these photoelectric standards leads to residuals of -0.036 and -0.069 in V and -0.055 and -0.022 in B. Considering the fact that on the one hand the photometric errors for these stars are about 0.05 magnitudes, and on the other hand some of our fitted standards have residuals up to 0.07 magnitudes, these residuals are in good agreement with the previous results.

We point out that the checks performed so far give us confidence that to date the relative photometry of the pipeline permits us to reduce and to co-add images to a photometric accuracy of about 0.05 magnitudes (the average between the standard deviations of the fits to secondary photometric standards of Pietrzynski et al., 2002, and photoelectric standard stars from Graham, 1981). This is in agreement with the expected errors due to the above mentioned illumination issues.

Table 3. Standard deviation of the global solution applied to each chip. The relation between chip number and actual position on the mosaic can be seen in Fig. 1.

Chip No.	$\sigma(B)$	$\sigma(V)$
1	0.035	0.024
2	0.040	0.034
3	0.039	0.033
4	0.040	0.031
5	0.030	0.026
6	0.036	0.033
7	0.039	0.033
8	0.039	0.027
1-8	0.037	0.028

5.6. *Set* quality characteristics

The object catalogues generated for the astrometric and photometric calibration give us a good overview of the data quality at this stage. They allow us to investigate in more detail the PSF properties of individual exposures which is crucial in weak lensing studies. Also the night sky conditions (transparency, sky brightness) can be studied and images that should be excluded from the final co-addition can be identified. Fig. 22 discusses PSF properties of WFI@2.2m and Fig. 23 shows an overview of our currently implemented quality control on the *Set* basis.

6. Weighting and Flagging

With astrometric and photometric information for all images at hand we can finally co-add them and produce a deep mosaic. This section motivates our choices for the implemented co-addition procedure which is very close to that of the EIS Wide Survey described in Nonino et al. (1999).

As long-exposed SCIENCE frames are dominated by Gaussian sky noise the optimal result (in terms of noise properties in the co-added image) is obtained by a weighted mean co-addition of all input images. However, this is not straightforward as many image pixels carry non-Gaussian noise properties (such as bad columns, vignetted image regions, cosmics) or other defects that we would like to exclude from the co-addition process (such as satellite tracks or extended stellar reflections). To perform a weighted mean co-addition all *bad* image pixels have to be known beforehand and assigned a zero weight in the co-addition process. We note that identifying defects before the co-addition has several advantages over rejecting them during the stacking process:

– Rejection methods during the co-addition, such as median stacking or sigma-clip rejection require a minimum number of 5-10 input images for robust statistics. With prior knowledge of the defects good results can already be obtained with a small number of images.

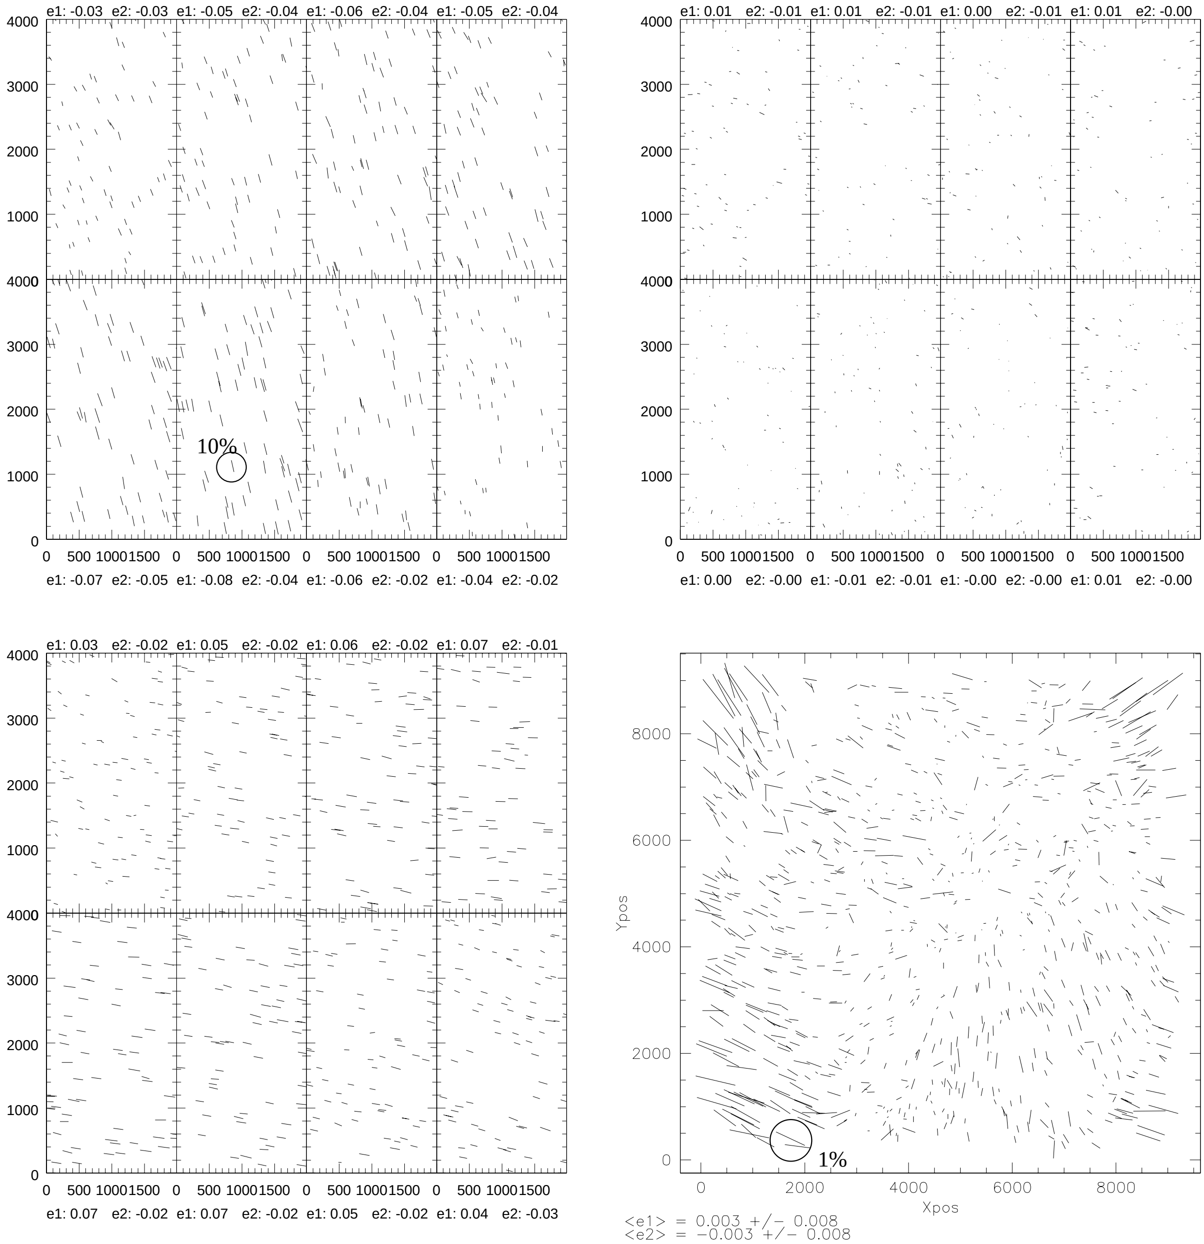

Fig. 22. PSF properties of WFI@2.2m fields: For each image within a *Set* we characterise PSF properties (PSF size and anisotropy) over the complete field-of-view of a WFI. To this end we run the KSB algorithm [a weak lensing analysis technique for accurate shape measurements of astronomical sources; see Kaiser et al. (1995)] over selected stars of all images. The upper left panel shows PSF anisotropies for an intra-focal, the upper right for a focal and the lower left for an extra-focal exposure. The sticks represent amplitude (given as $\sqrt{e_1^2 + e_2^2}$) and orientation $\phi = 0.5 \arctan(e2/e1)$ of the PSF anisotropy. The chosen scale for the stick length is the same for these three plots in order to show the increase in the anisotropies with respect to the focused exposure. The mean stellar ellipticities are 0.066, 0.09 and 0.059, respectively. The lower right panel depicts typical PSF anisotropies of a WFI@2.2m *R*-band mosaic ($\sim$ 57 exposures with $\sim$ 500 s exposure time each). Note that the largest PSF anisotropies in the mosaic are as small as ≈ 0.01. Compared to the other three PSF plots, a different scale for the stick length was used in order to clearly show the anisotropies. We note that WFI@2.2m has excellent PSF properties if data are obtained under favourable focus conditions.

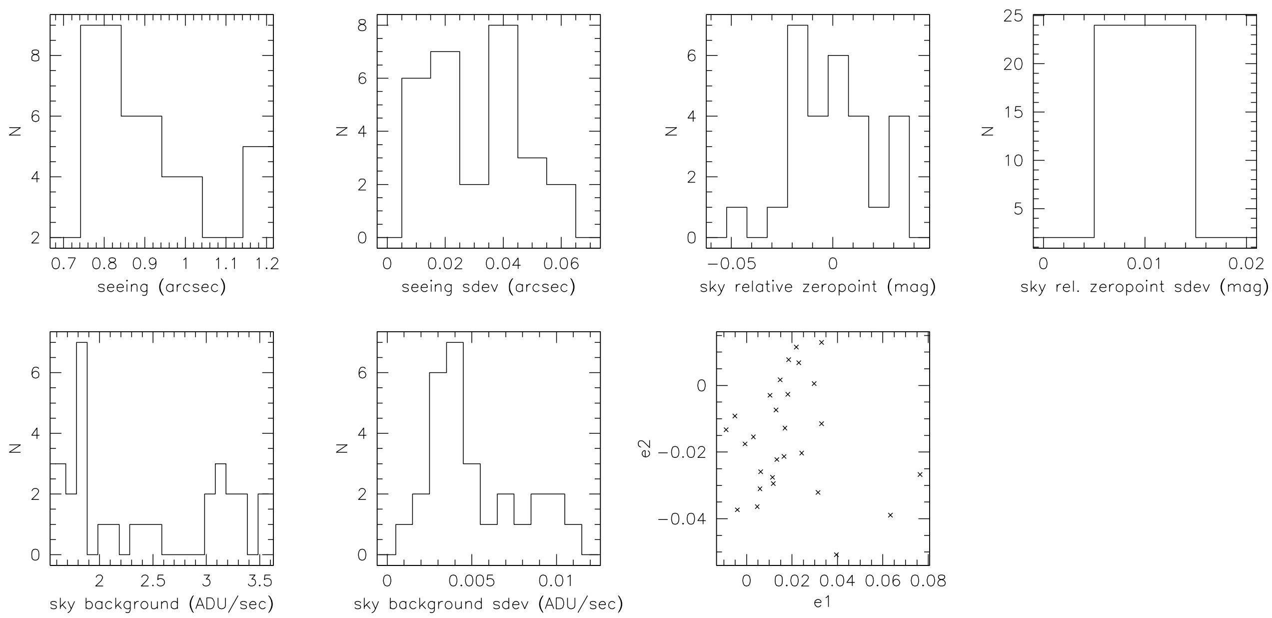

Fig. 23. Shown are quality control plots for a *Set* of 28 *R* band exposures from WFI@2.2m. They primarily allow us to select the images that enter the final co-addition process. The two upper left plots show the seeing distribution in the *Set*, where the seeing value of an exposure is estimated from the mean seeing in the individual chips. The distribution of seeing variations within an exposure is also shown (second left plot). The two lower left plots show the night sky amplitude (normalised to an exposure time of one second) during the observations. They allow us to identify exposures taken during twilight or unfavourable moon phases. In the upper right we have a quick check on the photometric conditions through the distribution of relative zero points (see Sect. 5.3). The standard deviations of relative zero points show that our sky-background equalisation (see Sect. 4.6) adjusts zero points typically to 0.01-0.03 mag within a WFI@2.2m mosaic. In the lower right we see the PSF ellipticity distribution (see Fig. 22). Here each point represents the mean ellipticity value from all stars within an exposure.

– Co-additions based on median or sigma-clip rejection are problematic when observations under varying seeing conditions have to be combined. In this case object profiles in the co-added image might be affected by the rejection algorithm.
– A practical advantage is that a weighted mean without rejections is a strictly linear process. Hence, new images can directly be added to an already existing mosaic and not all input images need to be co-added again.

We identify bad pixels in the individual exposures as follows:

1. Every chip contains *hot* and *cold* pixels, i. e. pixels that always have a high/low charge value. They are typically stable over the period of an observing run and are most effectively identified in the masDARK frames (see Fig. 24). If DARK frames are not at hand, SKYFLATs and/or SUPERFLATs often provide a sufficient bad pixel map.
2. Saturated image pixels are identified by applying a pixel value threshold to the SCIENCE frames.
3. Pixels affected by cosmic rays are not identified as easily as their appearance is different in each SCIENCE frame. We detect them with `SExtractor` with a special filter for cosmic ray detection generated with `EYE` (Bertin 2001; Nonino et al. 1999). The filter we actually use was specifically designed for cosmic ray removal on EMMI@NTT data in the framework of the EIS Wide project. It is difficult to say for which other instruments it can be used without modification or to quantify its actual performance. The visual impression from our co-added WFI@2.2m frames is that the detection efficiency for small-scale cosmic ray features is very high and gives a sufficient image cleaning for this camera (see Fig. 25).
4. Other extended defects such as satellite tracks or reflections from bright stars are masked by hand. They have been identified in the manual pass through the data during the *Run* processing (see Sect. 4.10).

In this way we can generate a FLAG map for each image. It is an integer FITS image containing zeros for every good pixel and values characterising the different defects otherwise. We use these FLAG maps in the creation of WEIGHT maps (see below) and to mark standard stars whose magnitude measurement is affected by a defect (see Sect. 4.11).[20]

Besides to the actual image co-addition algorithm we need to pay special attention to the noise properties in

[20] The FLAG and WEIGHT maps described in this section are actually created on the *Run* level before the standard field processing. Because of their very close relation to the image co-addition we shifted their description to the *Set* processing.

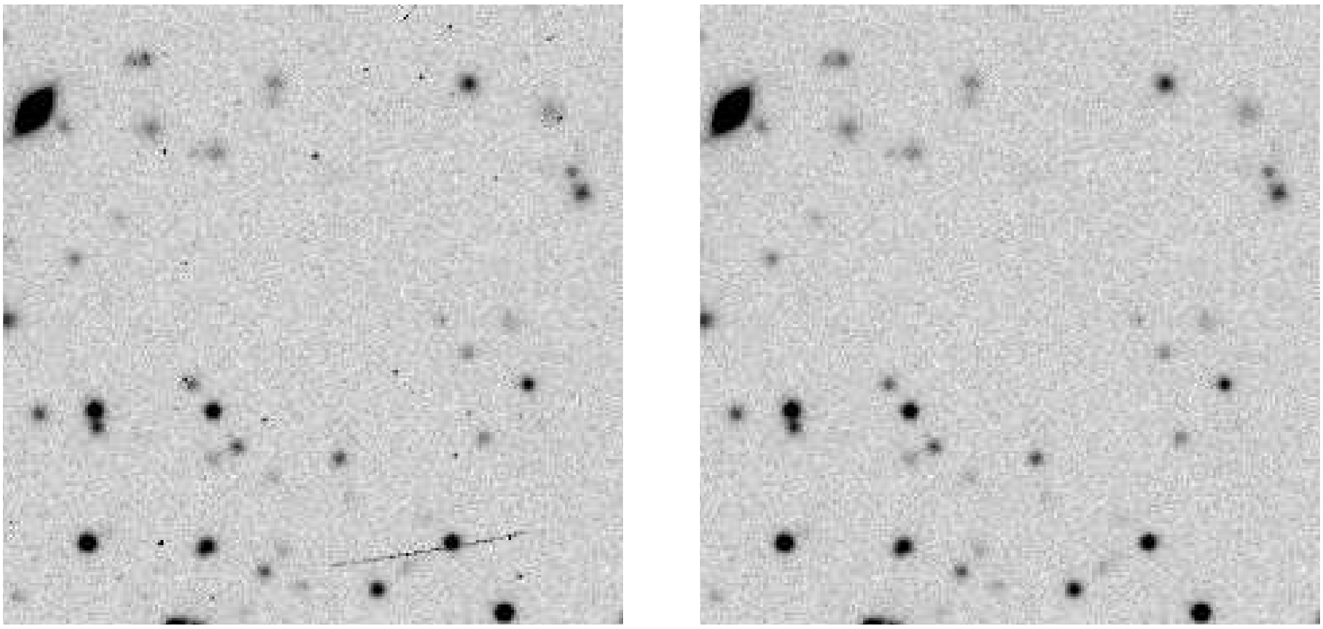

Fig. 25. The left panel shows 25 co-added WFI@2.2m frames where no cosmic ray cleaning has been performed. The right panel shows the result with the cleaning procedure turned on. We find that for WFI@2.2m frames `SExtractor` (in connection with a neural network filter generated with `EYE`) performs a very efficient detection of small-scale cosmic ray features.

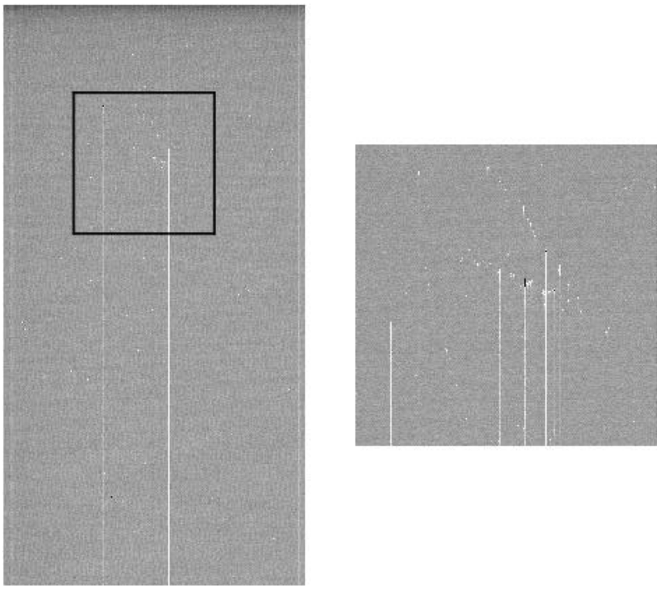

Fig. 24. Shown is a masDARK image and a zoomed in region. In a stack of several long-exposed DARK frames hot and cold pixels, which often come in groups and affect complete rows/columns, show up with high significance and are easily identified by applying a pixel value threshold to the masDARK.

the final co-added image. The gaps in multi-chip cameras and the different intrinsic gains of the various detectors lead to complex noise properties in a co-added image consisting of dithered exposures. These noise variations need to be taken into account properly when estimating a threshold for object detection or for estimating errors based on pixel noise (such as object fluxes for instance). For a full characterisation of the noise, another image besides the final co-added image, the WEIGHT map, describing the relative noise properties for each image pixel, is created during the co-addition process. We arrive at this WEIGHT map in the following way:

1. We produce a WEIGHT map for every input image to be co-added. The starting point for these individual WEIGHTs is the masFLAT that is rescaled to a mode of unity. It provides information on the relative sensitivity (and hence noise) variations within an image. This basic map is modified by the FLAG map by setting the value of defect pixels to zero.
2. These individual WEIGHTS are co-added together with the science images. The co-added WEIGHT provides a full characterisation of the relative noise properties later. The exact co-addition procedure is described in Sect. 7. Fig. 26 shows an example for a co-added WEIGHT map and Fig. 27 its merits for object detection.

FLAG and WEIGHT maps are produced with the TERAPIX tool `weightwatcher` (see Bertin 1998).

7. Image co-addition

Before the final co-addition, all images are brought to the same background level by subtracting the night sky. To accurately model the sky-background and to avoid biased estimates close to large astronomical sources we apply the following two-stage process:

1. We run `SExtractor` and detect all large-scale objects having at least 50 adjacent pixels with 1.5σ over the sky-background (these values are for WFI@2.2m and can vary according to the data set). All image pixels

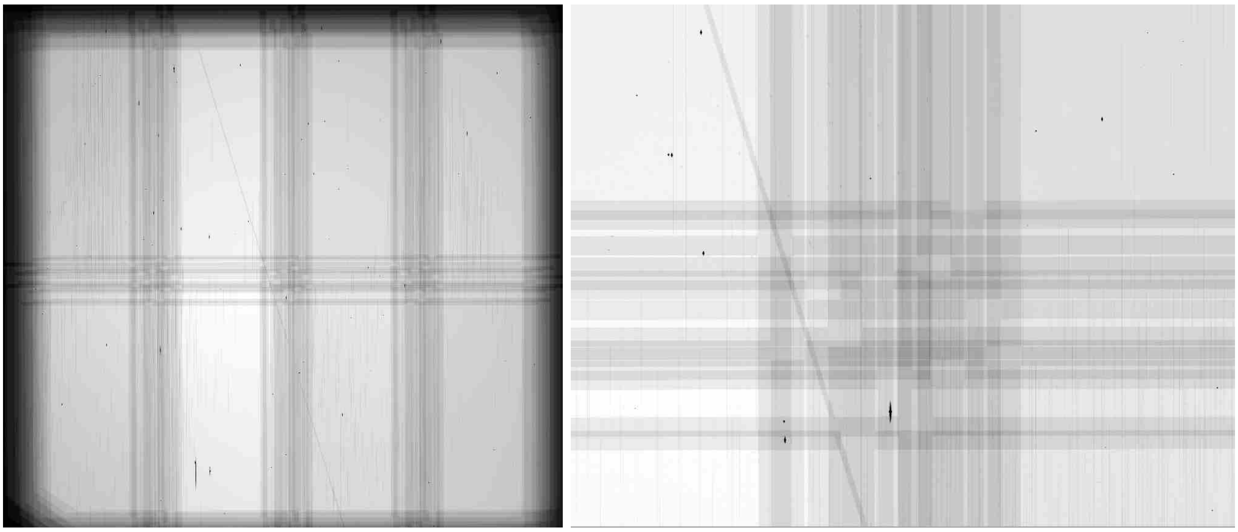

Fig. 26. Shown is the WEIGHT map (left panel) characterising the noise properties of a co-added science image and a zoom-in to its centre (right panel). The lighter the colour, the higher the relative weight (the lower the noise) of the pixel. The darker regions at the positions where different chips meet have about half the weight of well-covered regions. Different weights between regions where the same number of input images have contributed show intrinsic sensitivity variations. We note that the noise structure is quite complex and cannot be taken into account appropriately without the WEIGHT map. These maps are used by `SExtractor` in the object detection phase and for noise calculations. Having the map at hand one also does not need to cut the outermost regions of the coadded images where the noise is considerably higher than in the inner regions. See also Fig. 27.

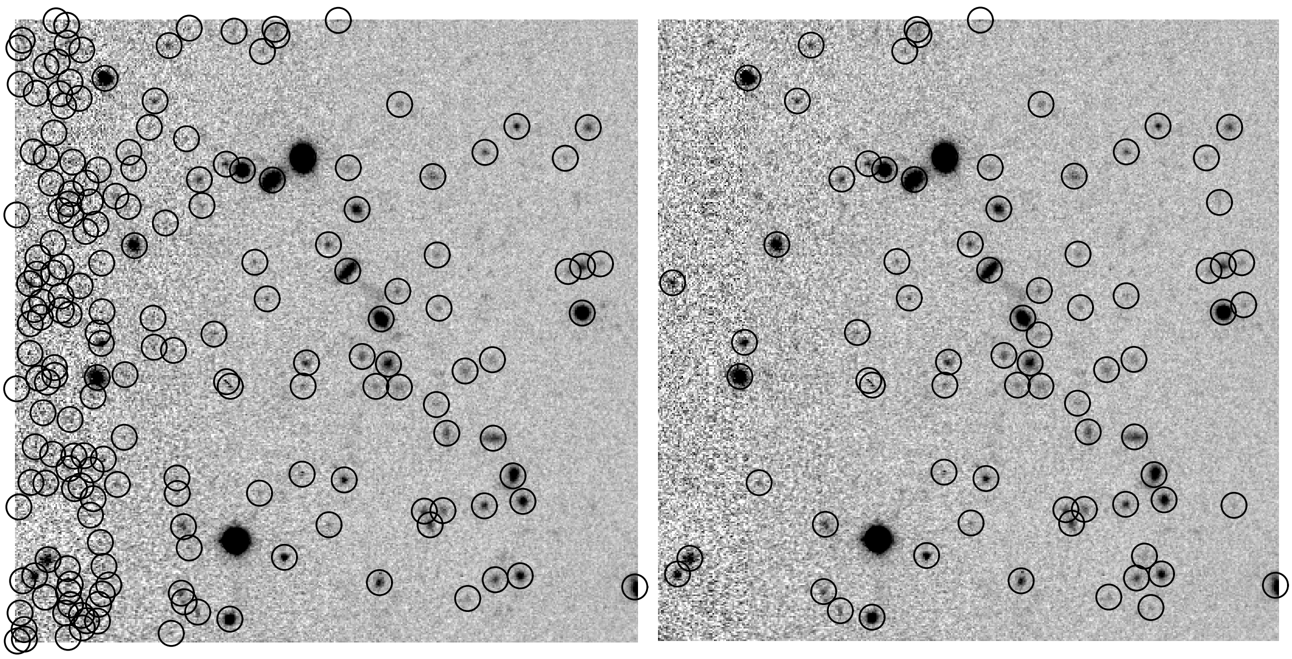

Fig. 27. `SExtractor` source catalogue obtained from the same mosaic (with an area of $95'' \times 95''$) without (left) and with (right) a WEIGHT image as an additional `SExtractor` argument. The number of spurious detections in regions with higher noise is obvious when no WEIGHT image is used as an additional input.

belonging to detected objects are replaced with the image mode.

2. From this object-subtracted image we create a `SExtractor` BACKGROUND check-image (BACK_SIZE=90 for WFI@2.2m images) and subtract it from the original SCIENCE image.

We showed in Sect. 5.4 that this sky subtraction does not introduce a significant bias in the object photometry on the co-added images.

For the co-addition we can use Richard Hook's `EIS Drizzle` which is implemented within IRAF or `SWarp` (E. Bertin, TERAPIX). Both software packages perform a weighted mean co-addition on a sub-pixel level taking

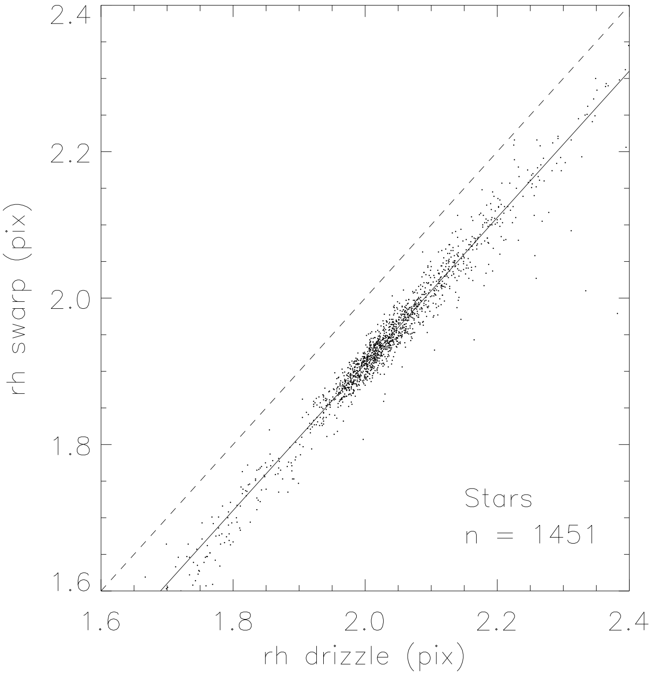

Fig. 28. Comparison of the image seeing between two mosaics created with `SWarp` and `EIS Drizzle`. Shown are the half-light radii for unsaturated stars in the two co-additions. The swarped image has an image seeing that is 0.09 pixels smaller than the one for the drizzled image. The mean values for the image seeing are 0.″95 for the drizzled image and 0.″91 for the swarped one. The same astrometric solution was used for both co-additions. Thus `EIS Drizzle` slightly increases the size of the PSF, an effect of its simplified kernel. In Fig. 30 it is shown that the PSF anisotropies are identical for both co-addition strategies.

into account the full astrometric solution and allowing a variable output pixel size. `EIS Drizzle` was developed within the EIS Wide project (Nonino et al. 1999) to co-add quickly large volumes of data with low demands on CPU and virtual memory. It uses strongly simplified calculations during the pixel resampling process and allows only strictly linear co-additions (see below). In contrast, `SWarp` offers a large variety of pixel resampling and co-addition algorithms.

`SWarp` first undistorts and resamples all input SCIENCE and WEIGHT images according to the astrometric solution. The user can choose between several sophisticated kernels for the pixel remapping. We use by default the LANCZOS3 kernel (see Bertin 2002, for details on `SWarp's` resampling kernels). The final co-addition of the resampled images is performed in a second pass. Having all resampled input images belonging to a given output pixel at hand simultaneously, `SWarp` can calculate the final result in a variety of ways, such as median, mean or the weighted mean which is our method of choice. `EIS Drizzle's` implementation is significantly different. It uses a method known as forward mapping. It performs the complete co-addition in one pass by putting the input images consecutively on the output grid. Hence, `EIS Drizzle` is limited to linear co-addition methods. Applying the astrometric solution, an individual input pixel is mapped somewhere on top of several neighbouring output pixels, and is distorted and rotated. Its flux is then distributed accordingly amongst these output pixels.[21] The `EIS Drizzle` approach strongly simplifies the calculation of the flux distribution, in the sense that only non-integer shifts are taken into account, whereas rotations and distortions of the mapped pixels are neglected. This corresponds to the *Turbo Kernel* in the current `MultiDrizzle` implementation (Koekemoer et al. 2002). See also Fruchter & Hook (2002) for a more detailed description of the drizzling approach. We discuss the differences of the pixel resampling kernels of `SWarp` and `EIS Drizzle` in more detail at the end of this section.

Choosing a weighted mean co-addition for `SWarp`, both software packages calculate the value of an output pixel in the same way. Four factors contribute to the final result in our weighted mean co-addition. Given are the value I_i, representing the part of an input pixel that goes to a specified output pixel $I_{\rm out}$ in the co-added image. W_i represents the associated value in the WEIGHT map. I_i is scaled with factors f_i to the consistent photometric zero point and normalised to a fixed exposure time of 1 s. This scaling reads

$$f_i = 10^{-0.4\,ZP_i}/t_i, \tag{6}$$

where t_i is the exposure time and ZP_i the relative photometric zero point. All images are in addition weighted according to their sky noise. This weight scale is given by

$$w_i = \frac{1}{\sigma^2_{{\rm sky},i} f_i^2}. \tag{7}$$

We take into account that the noise also scales with the flux scale f_i. The values $I_{\rm out}$ and $W_{\rm out}$ in a stack of N exposures then read

$$I_{\rm out} = \frac{\sum_{i=1}^N I_i f_i W_i w_i}{\sum_{i=1}^N W_i w_i}\;, \qquad W_{\rm out} = \sum_{i=1}^N W_i w_i\;. \tag{8}$$

Besides the co-added SCIENCE and WEIGHT mosaics `EIS Drizzle` can produce a CONTEXT map which allows the identification of all input frames that contribute to a given pixel in the co-added images.

As standard sky projection of the mosaic we use the TAN projection (see Greisen & Calabretta 2002, for further information on sky projections) with an orientation such that North is up and East to the left. A reference coordinate can be specified for the co-addition, associated with an integer reference pixel. Thus, if multi-colour information is available for a particular *Set*, the mosaics in the different bands can be registered with sub-pixel accuracy if required.

In the following we show a comparison of photometric properties from objects extracted from the same *Set* co-added with `EIS Drizzle` and `SWarp`. We expect that the

[21] The user can vary the fraction of the input pixel entering this procedure (drizzle PIXFRAC parameter). We always use the original pixel size in our co-additions (PIXFRAC=1).

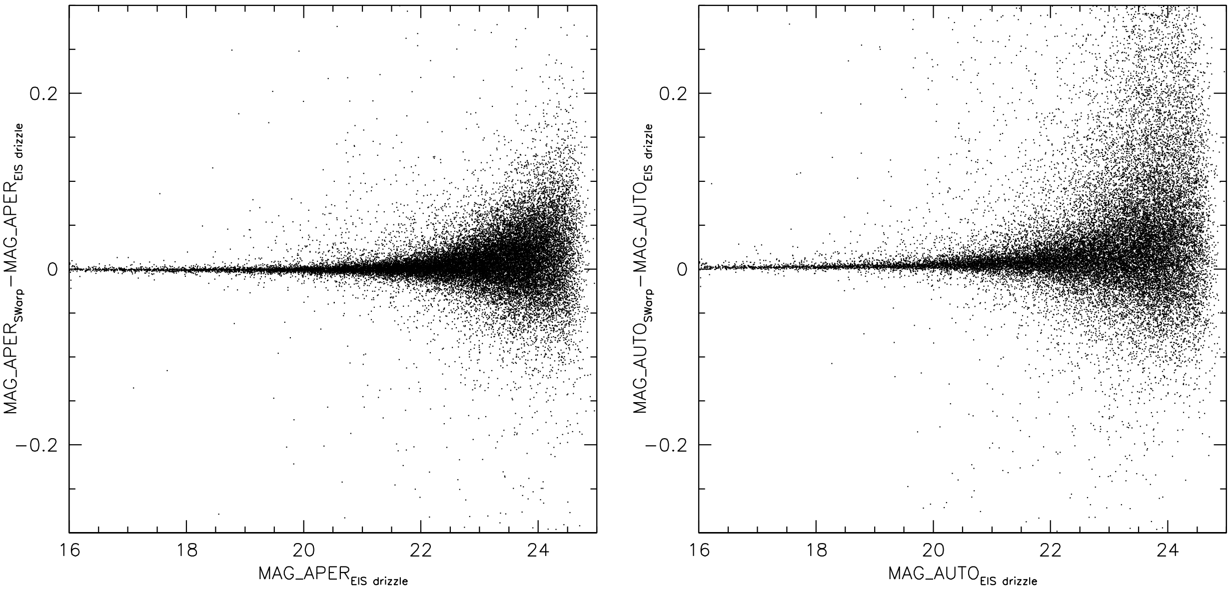

Fig. 29. We show comparisons of magnitude estimates from an `EIS Drizzle` and a `SWarp` co-addition from the same *Set*. Both measurements are in very good agreement if the flux is measured within a fixed aperture (here with a diameter of 3.″5). However, for `SExtractor's` MAG_AUTO estimate (Kron like total magnitudes), the aperture used for faint sources is larger in the drizzled images than in the `SWarp` mosaic. The reason is the more correlated noise in the `EIS Drizzle` approach.

simpler resampling kernel from `EIS Drizzle` leads to stronger noise correlations and image blurring than the more sophisticated `SWarp` approach. This is illustrated in Fig. 28 and Fig. 29. The `EIS Drizzle` co-addition has a slightly larger image seeing and `SExtractor` chooses a larger *optimal* radius for its MAG_AUTO parameter. In contrast, `EIS Drizzle` has no significant impact on the PSF anisotropies (see Fig. 30). Also more sophisticated weak lensing analyses (such as cluster mass reconstructions) revealed that results based on object shape measurements do not differ significantly between the two co-additions.

8. First Data quality assessments of co-added *Sets*

The implementation of a thorough, automatic quality check on our co-added mosaics is still in its infancy. In this section we describe basic tests on extracted object catalogues.

8.1. Galaxy counts

The programme `SExtractor` is used to create a raw catalogue of all objects that consist of at least 3 contiguous pixels (DETECT_MINAREA=3) with a flux 2σ above the flux of the sky-background (DETECT_THRESH=2). The source extraction is done on a filtered image; we use a normalised Gaussian filter with a full width half maximum of 4.0 pixels (FILTER_NAME=gauss_4.0_7x7.conv). This conservative threshold is chosen in order to minimise the number of spurious detections. In the following, we use the `SExtractor` parameters MAG_AUTO for magnitudes and FLUX_RADIUS for the half-light radius. To create a galaxy catalogue, the `SExtractor` parameter CLASS_STAR in combination with the half-light radius is used; saturated objects are rejected. We define every object which has CLASS_STAR less than 0.95 as a galaxy. A simple check of this selection is a magnitude over half-light radius plot. All stars have the same half-light radius and therefore show up as a vertical branch in this plot, see Fig. 31.

We count the number of galaxies in 0.5 mag wide bins per one square degree. To normalise the area to one square degree, we take into account that each object occupies an area in which fainter objects can not be detected. For this correction we use the `SExtractor` parameter ISO0, which is the isophotal area above the analysis threshold. We note that in the case of empty fields this effect is almost negligible. An error-weighted linear regression to the logarithmic galaxy counts is performed and the slope, $\mathrm{d}\log N/\,\mathrm{d\,mag}$ determined. We routinely compare our galaxy counts with those of McCracken et al. (2003), see Fig. 32. With this comparison, a rough test of the magnitude zero point and the limiting magnitude can be performed.

8.2. Clustering of extended sources

A further test for the quality of the co-added image is the clustering of sources. For that purpose we use the two-point angular correlation function, $\omega(\theta)$, where $\omega(\theta)\delta\theta$ is the excess probability of finding a pair separated by an angle between θ and $\theta + \delta\theta$. We estimate this quantity by creating a large number of random catalogues (by

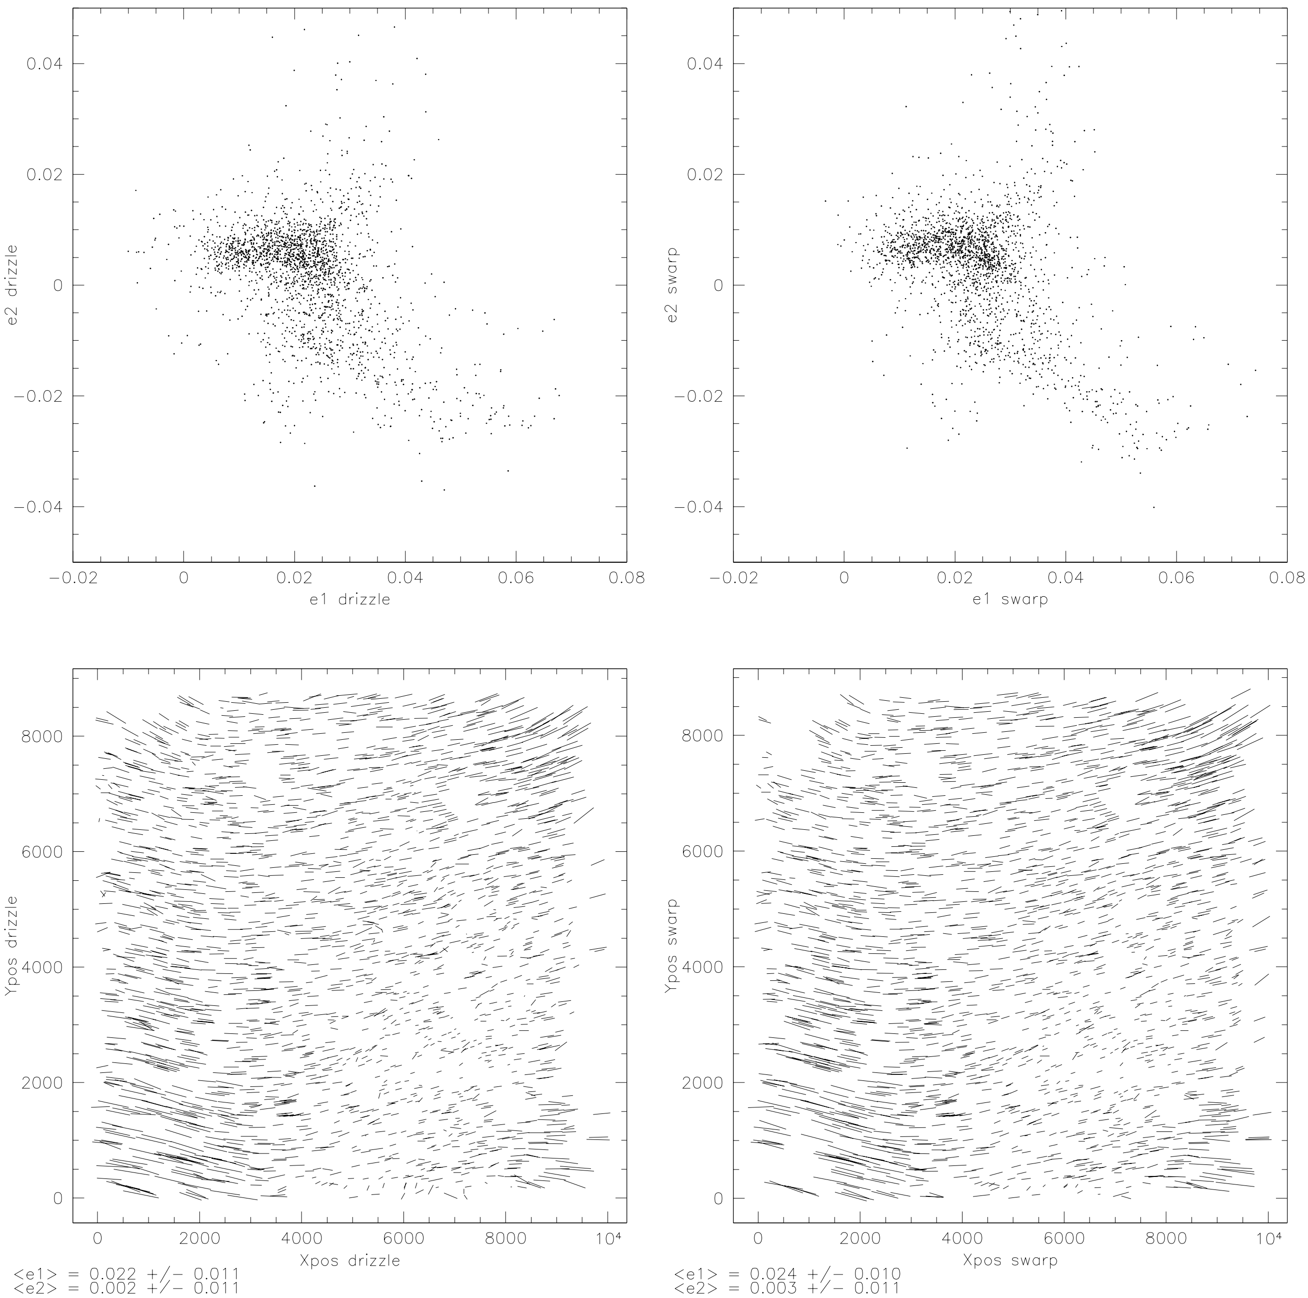

Fig. 30. PSF anisotropies in a drizzled (left column) and a swarped (right column) mosaic of the same data *Set*. The patterns are virtually identical. Yet `EIS Drizzle` marginally increases the size of the PSF, as was shown in the left panel of Fig. 28. The mean PSF anisotropy for this particular mosaic amounts to mere 0.022 ± 0.015.

default 40 mock catalogues are created) and count the pairs within the data catalogue, DD, within the random catalogue, RR and between the data and random catalogues, DR. The estimator for $\omega(\theta)$, proposed by Landy & Szalay (1993), is

$$\omega(\theta) = \frac{DD - 2DR + RR}{RR}. \tag{9}$$

The random fields must have the same geometry as the data field. Therefore we calculate an obscuration mask out of the number density of extracted sources as follows. A mesh with 512×512 mesh cells is placed on top of a data field and the number of objects in each cell is counted. Then the cell count matrix is smoothed with a Gaussian kernel of about 512/60 cells FWHM, and all matrix elements lower than a given fraction (we use 75% as default but varying this parameter between 50% and 80% does not change the results significantly) of the mean number of galaxies inside a cell are defined as a masked region. The borders of a field, bright stars, and part of their halos are masked automatically by this method (see Fig. 33).

To maximise computational speed, we perform our calculations by creating an index tree for galaxy position as explained in Zhang & Pen (2004). The error bars in

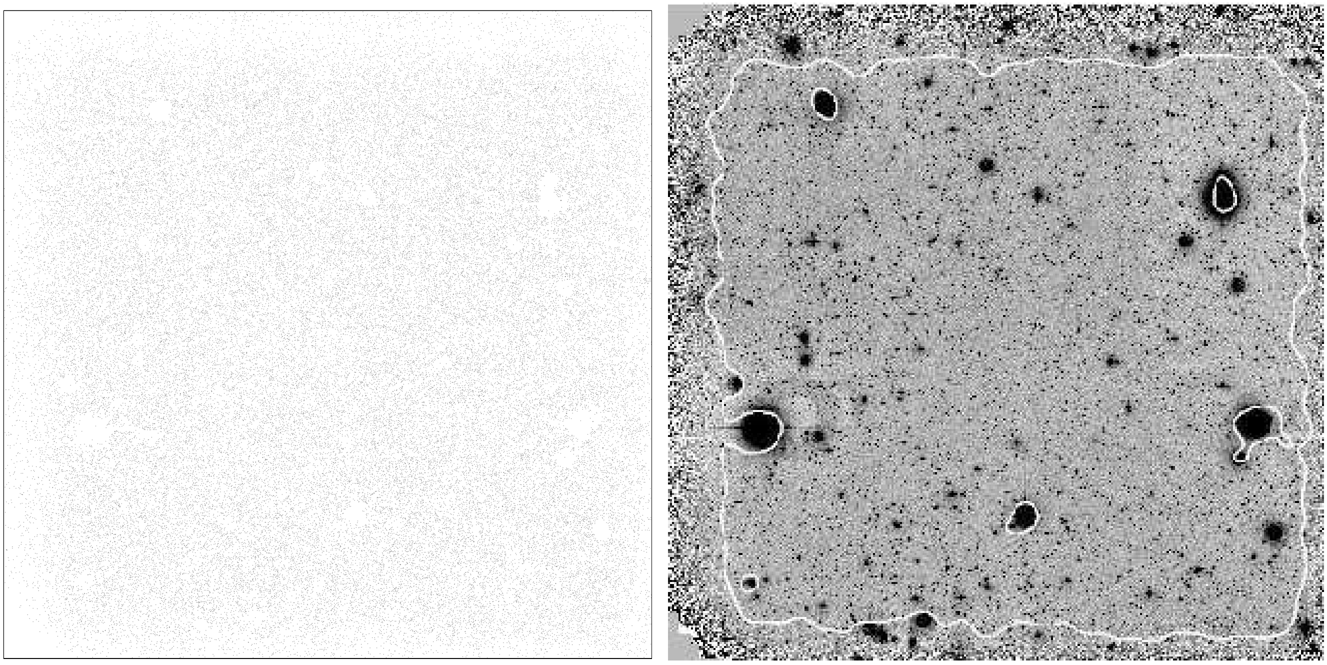

Fig. 33. Obscuration mask for the CDFS. Left panel: Distribution of detected sources. Right panel: FITS image of the CDFS. White contours indicate masked region. In deep fields, our automatic obscuration mask based on object density variations reliably marks large-scale astronomical sources and noisy borders which would significantly influence the area of our correlation function analysis. This approach turned out to be sufficient for quality check purposes.

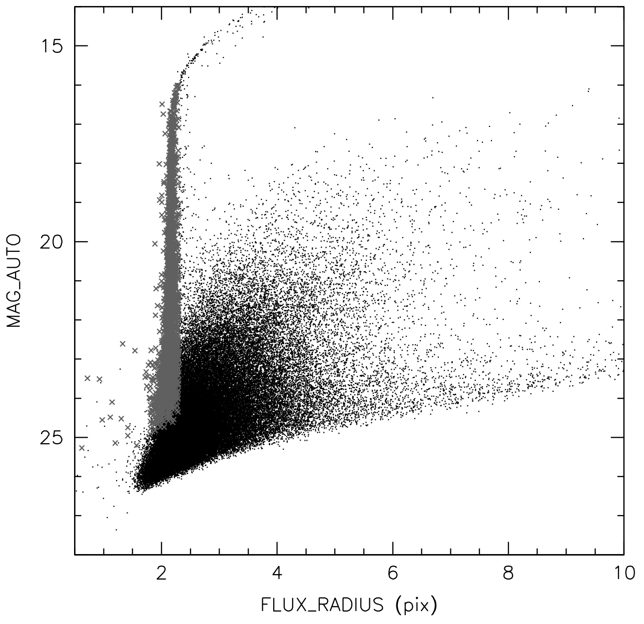

Fig. 31. Magnitude vs. half-light radius plot. Grey crosses: stars; appear as a vertical branch in the plot and are selected by the CLASS_STAR parameter and half-light radius. Dots: extended sources.

our check-plots for each angular bin are simply estimated by Poisson noise and are therefore a lower limit to the uncertainty in $\omega(\theta)$.

As an example we present the two-point angular correlation function of galaxies in our R-band reductions for 11 WFI-fields in the EIS Deep Public Survey (EIS DPS). For three different magnitude bins we fitted a power law, $\omega(\theta) = A\,\theta^{\delta} - C$, to the data. The variable C is the so-called integral constraint (see for instance Roche et al. 1993), which only becomes important for large fields at larger scales ($\theta > 2'$). We therefore neglect the integral constraint and perform the fit for small scales ($\theta < 2'$). The results are shown in Fig. 34.

The results from our clustering analysis can be cross-checked by considering the aperture number count dispersion $\langle N^2\rangle(\theta)$. It is directly related to the angular correlation function $\omega(\theta)$ by

$$\langle N^2\rangle(\theta) = \int \mathrm{d}\vartheta\, \frac{\vartheta}{\theta^2}\, \omega(\vartheta)\, T_+\left(\frac{\vartheta}{\theta^2}\right), \tag{10}$$

where the function T_+ is defined as eq. (35) in Schneider et al. (2002) and it has the nice property of being independent of the integral constraint. As an example we use the measured $\omega(\theta)$ for the magnitude interval $R \in [19.5; 21]$ from the 11 WFI-fields of the last section to calculate $\langle N^2_{\rm obs}\rangle(\theta)$, see Fig. 35. To compare the slope δ and amplitude A_1 of the fitted power law to the measured $\omega(\theta)$ of the previous section (Fig. 34), $\langle N^2_{\rm fit}\rangle(\theta)$ is calculated for the power law, $\omega(\theta) = A_1\,\theta^{\delta}$. The results show that the slope δ obtained from the power law fit of the angular correlation function is correct. The amplitude A_1, however, seems, with $A_1 = 0.133$, to be a bit larger; this small discrepancy can easily arise due to the fact that $\langle N^2_{\rm obs}\rangle(\theta)$ is calculated using $\omega(\theta)$ over the entire θ-range and that the ω-fit is slightly influenced by the integral constraint. As a comparison, Fig. 35 displays the function $\langle N^2\rangle(\theta)$ for the fit parameters determined by Hoekstra et al. (2002) for $R_{\rm C}$-band data of the same

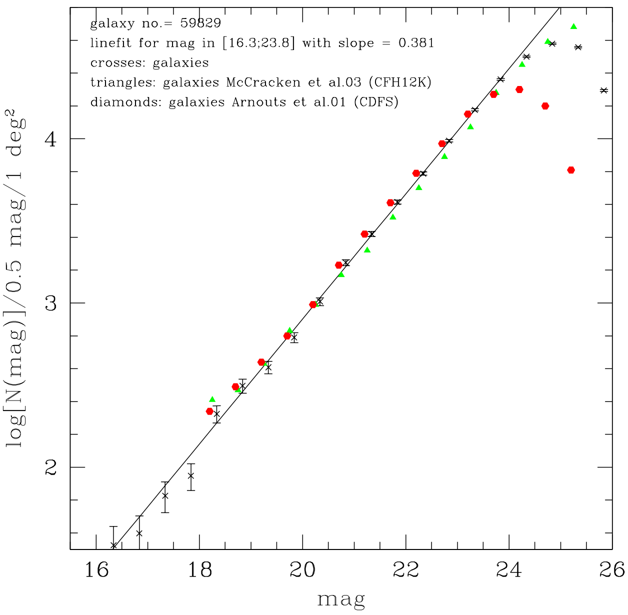

Fig. 32. The figure displays the logarithmic galaxy counts in 0.5 mag bins per one square degree from our reduction of the Chandra Deep Field South (CDFS; see Giavalisco et al. 2004). The error bars are due to Poisson noise; the line fit is an error-weighted linear regression in a magnitude range between the saturation and the limiting magnitude (here: $R \in [16.3; 23.8]$). For the normalisation of the area we take into account that each detected object occupies an area in which fainter objects cannot be detected. As a comparison to our galaxy counts we also plot the galaxy number counts from the CFH12K-VIRMOS deep field (McCracken et al. 2003). For the CDFS we also plot the number counts from Arnouts et al. (2001).

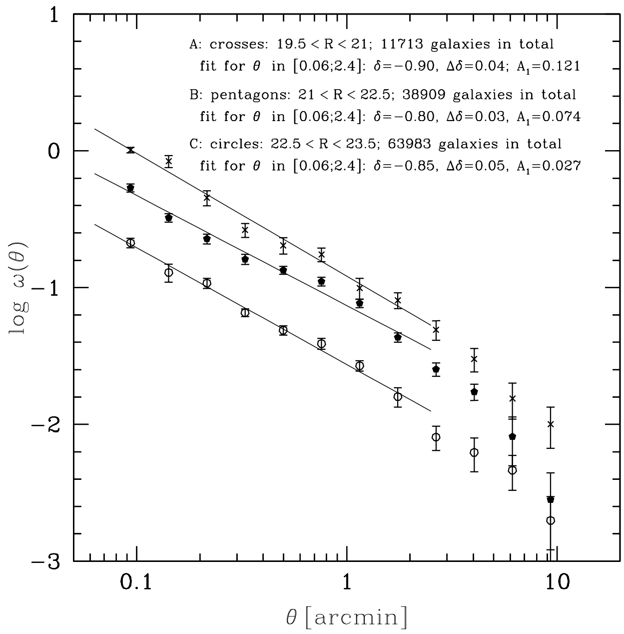

Fig. 34. The two-point angular correlation function of galaxies for different magnitude bins, for 11 WFI-fields in the DPS. We show the logarithm of the amplitude, log $\omega(\theta)$, as a function of the angular separation θ in arcmin. Here the error bar in each bin is due to the field-to-field variance. We perform a simple error-weighted linear regression in the angular interval $\theta \in [0\farcm06; 2']$ (where the integral constraint is negligible) to determine the slope δ and the amplitude $A_1 = \omega(\theta = 1')$.

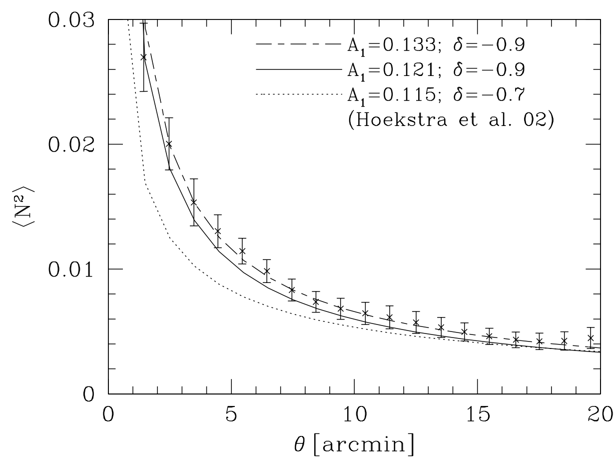

Fig. 35. The aperture number count dispersion $\langle N^2_{\rm obs}\rangle(\theta)$ as a function of angular scale θ in arcmin for 11 WFI-fields in the magnitude interval $R \in [19.5; 21]$. The error bars on the function $\langle N^2\rangle(\theta)$ are due to the field-to-field variance of the 11 fields. Note that the points are correlated. The lines display the function $\langle N^2_{\rm fit}\rangle(\theta)$ assuming a power law for the angular two-point correlation function of the form $\omega(\theta) = A_1\,\theta^\delta$ for different parameters $A_1 = \omega(\theta = 1')$ and δ. Solid line: $\langle N^2_{\rm fit}\rangle(\theta)$ calculated for the fit parameter obtained from the angular correlation function (see Fig. 34); dotted line: comparison with $\langle N^2\rangle(\theta)$ calculated for the fit parameters determined by Hoekstra et al. (2002); dashed line: best fit to $\langle N^2\rangle(\theta)$.

magnitude range. Note that this can only be a rough comparison, because of different source extraction algorithms and slightly different R-band filter.

9. Conclusions and Outlook

We have presented our image processing methods for multi-chip cameras that we developed in the course of the GaBoDS project. A significant fraction of GaBoDS is a virtual survey in which observational data were collected from the ESO Science archive within an ASTROVIRTEL program[22] primarily for weak gravitational lensing studies. This allowed us to test and to apply our procedures on data sets which were acquired for a large variety of scientific programmes (e.g. deep field observations, search for moving objects, Cepheid studies in nearby galaxies) and obtained in many different ways (very compact to very wide dither patterns, *Sets* observed within a single night or over several years). Our experiences regarding the processing of these data sets from WFI@2.2m are the following:

[22] ASTROVIRTEL Cycle 2: Gravitational lensing studies in randomly distributed, high galactic latitude fields; P.I. Erben

– The techniques described perform very well on empty field observations with WFI@2.2m. The excellent optics of this instrument allow an accurate astrometric alignment of images which is crucial for weak lensing studies (Schirmer et al. 2003, 2004). Also very large data *Sets*, obtained over several years, can be processed. For example, we collected and reduced all WFI@2.2m observations from the CDFS, which consist of more than 100 individual exposures in each of the U, B, V and R bands (Giavalisco et al. 2004; Mobasher et al. 2004).
– Data *Sets* from crowded fields or from large-scale objects, whose extent is comparable to the field-of-view, can be processed with good results but require substantial manual intervention (see Schirmer et al. 2003, for processing details of the field around NGC 300).
– Our absolute photometric calibration is currently accurate to about 0.05 mag as discussed in Sect. 5.3 and Sect. 5.5. Also the comparison of our standard star calibrations with independent measurements in Sect. 4.11 shows errors of the same order. However, it seems that non-uniform illuminations that probably contribute most to the error budget, do not strongly depend on wavelength and hence errors on colours are smaller (see Koch et al. 2004). Our reduction of the CDFS data gives good results in photometric redshift studies of different groups (Mobasher et al. 2004; Gabasch et al. 2004; Hildebrandt et al. 2004).

As was described in Sect. 2 our processing system has already been used with a larger variety of single- and multi-chip cameras. However, more sophisticated tests on the astrometric and photometric accuracies, as presented here for WFI@2.2m, have not been done except for the single-chip camera FORS1@VLT (Bradač et al. 2004) and FORS2@VLT with a two-chip camera (Erben et al. 2003).

We are currently performing a thorough comparison of released fields from the EIS DPS with our own reductions of this survey. This will give us additional insights in the quality and the properties of our reductions and allows us to identify useful and necessary quality tests for the co-added mosaics. This work will be presented in Hildebrandt et al. (in prep.).

Acknowledgements. We thank Emmanuel Bertin and Oliver Czoske for comments and suggestions on the manuscript, and are deeply grateful to Yannick Mellier and Ludovic van Waerbeke for their support of the GaBoDS project and their long-standing collaboration. We are very thankful to Jean-Charles Cuillandre for making available an early version of his FLIPS software. The MPA and the MPE in Garching, the TERAPIX data centre at IAP in Paris and the Astrophysics department of the University of Innsbruck kindly gave us access to various computer platforms. We greatly appreciate the efforts of the people who helped us to improve our image processing system by using our tools, adjusting them to new instruments, making suggestions for their improvement and by reporting and/or fixing bugs. Many of us learned data reduction techniques while working within the ESO Imaging Survey Team at ESO.

This work was supported by the German Ministry for Science and Education (BMBF) through the DLR under the project 50 OR 0106 and through DESY under the project 05AE2PDA/8, and by the Deutsche Forschungsgemeinschaft (DFG) under the project SCHN 342/3–1. The support given by ASTROVIRTEL, a Project funded by the European Commission under FP5 Contract No. HPRI-CT-1999-00081, is acknowledged.

References

Alcalá, J. M., Radovich, M., Silvotti, R., et al. 2002, in Survey and Other Telescope Technologies and Discoveries. Edited by Tyson, J. A.; Wolff, S. Proceedings of the SPIE, Volume 4836., p. 406
Arnouts, S., Vandame, B., Benoist, C., et al. 2001, A&A, 379, 740
Bartelmann, M. & Schneider, P. 2001, Phys. Rep., 340, 291
Bertin, E. 1998, Weightwatcher v1.2 User's guide
—. 2001, Eye v1.1 User's guide
—. 2002, Swarp v1.34 User's guide
—. 2003, SExtractor v2.3 User's manual
Bertin, E. & Arnouts, S. 1996, A&AS, 117, 393
Bertin, E., Mellier, Y., Radovich, M., et al. 2002, in ASP Conf. Ser. 281: Astronomical Data Analysis Software and Systems XI, p. 228
Bradač, M., Erben, T., Schneider, P., et al. 2004, preprint astro-ph/0410642
Cuillandre, J., Luppino, G. A., Starr, B. M., & Isani, S. 2000, in Proc. SPIE Vol. 4008, Optical and IR Telescope Instrumentation and Detectors, Masanori Iye; Alan F. Moorwood; Eds., p. 1010
Devillard, N. 2001, in ASP Conf. Ser. 238: Astronomical Data Analysis Software and Systems X, p. 525
Erben, T., Miralles, J. M., Clowe, D., et al. 2003, A&A, 410, 45
Erben, T., van Waerbeke, L., Bertin, E., Mellier, Y., & Schneider, P. 2001, A&A, 366, 717
Fruchter, A. S. & Hook, R. N. 2002, PASP, 114, 144
Gabasch, A., Salvato, M., Saglia, R. P., et al. 2004, ApJ, 616, L83
Giavalisco, M., Ferguson, H. C., Koekemoer, A. M., et al. 2004, ApJ, 600, L93
Graham, J. A. 1981, PASP, 93, 29
Greisen, E. W. & Calabretta, M. R. 2002, A&A, 395, 1061
Hanisch, R. J., Farris, A., Greisen, E. W., et al. 2001, A&A, 376, 359
Hildebrandt, H., Bomans, D. J., Erben, T., et al. 2004, preprint astro-ph/0412375
Hoekstra, H., Van Waerbeke, L., & Gladders, M. D. 2002, ApJ, 577, 604
Kaiser, N., Squires, G., & Broadhurst, T. 1995, ApJ, 449, 460
Koch, A., Grebel, E., Odenkirchen, M., & Caldwell, J. A. R. 2004, AN, 325, 299

Koekemoer, A. M., Fruchter, A. S., Hook, R. N., & Hack, W. 2002, in The 2002 HST Calibration Workshop : Hubble after the Installation of the ACS and the NICMOS Cooling System, Proceedings of a Workshop held at the Space Telescope Science Institute, Baltimore, Maryland, October 17 and 18, 2002. Edited by Santiago Arribas, Anton Koekemoer, and Brad Whitmore. Baltimore, MD: Space Telescope Science Institute, 2002., p. 339

Koranyi, D. M., Kleyna, J., & Grogin, N. A. 1998, PASP, 110, 1464

Landolt, A. U. 1992, AJ, 104, 340

Landy, S. D. & Szalay, A. S. 1993, ApJ, 412, 64

Magnier, E. A. & Cuillandre, J.-C. 2004, PASP, 116, 449

Manfroid, J., Selman, F., & Jones, H. 2001, ESO Messenger, 102, 16

McCracken, H. J., Radovich, M., Bertin, E., et al. 2003, A&A, 410, 17

Mobasher, B., Idzi, R., Benítez, N., et al. 2004, ApJ, 600, L167

Monet, D. B. A., Canzian, B., Dahn, C., et al. 1998, VizieR Online Data Catalog, 1252

Nonino, M., Bertin, E., da Costa, L., et al. 1999, A&AS, 137, 51

Pietrzyński, G., Gieren, W., & Udalski, A. 2002, PASP, 114, 298

Radovich, M. 2002, ASTROMETRIX

Radovich, M., Arnaboldi, M., Ripepi, V., et al. 2004, A&A, 417, 51

Reif, K., Klink, G., Müller, P., & Poschmann, H. 2004, in Scientific Detectors for Astronomy, The Beginning of a New Era; eds. Amico, P.; Beletic, J. W.; Beletic, J. E., p. 367

Roche, N., Shanks, T., Metcalfe, N., & Fong, R. 1993, MNRAS, 263, 360

Schirmer, M., Erben, T., Schneider, P., et al. 2003, A&A, 407, 869

Schirmer, M., Erben, T., Schneider, P., Wolf, C., & Meisenheimer, K. 2004, A&A, 420, 75

Schneider, P., Van Waerbeke, L., & Mellier, Y. 2002, A&A, 389, 729

Stetson, P. B. 1987, PASP, 99, 191

—. 2000, PASP, 112, 925

Tyson, N. D. & Gal, R. R. 1993, AJ, 105, 1206

Valdes, F. G. 2002, in Automated Data Analysis in Astronomy, 309

Vandame, B. 2002, in Astronomical Data Analysis II. Edited by Starck, Jean-Luc; Murtagh, Fionn D. Proceedings of the SPIE, Volume 4847., p. 123

Vandame, B. 2004a, Alambic v1.0 User's guide

—. 2004b, Traitements d'images a grand-champs et multi-longeurs d'ondes (PhD Thesis)

Zacharias, N., Urban, S. E., Zacharias, M. I., et al. 2004, AJ, 127, 3043

Zhang, L. L. & Pen, U. 2004, preprint astro-ph/0305447

Appendix A: Pipeline Image header

Our pipeline replaces the original FITS headers of all individual CCDs by a new one containing only a minimum set of keywords. In this way we unify the headers for all instruments and avoid inconsistencies especially in the astrometric calibration. An example header for WFI@2.2m is:

```
SIMPLE  =                    T  /
BITPIX  =                   16  /
NAXIS   =                    2  /
NAXIS1  =                 2142  /
NAXIS2  =                 4128  /
BSCALE  =                   1.  /
BZERO   =               32768.  /
CTYPE1  =           'RA---TAN'  /
CTYPE2  =           'DEC--TAN'  /
CRPIX1  =                -416.  /
CRPIX2  =                -224.  /
CD1_1   =            -6.61E-05  /
CD2_2   =             6.61E-05  /
CD1_2   =                   0.  /
CD2_1   =                   0.  /
CRVAL1  =            12.505009  /
CRVAL2  =            -52.15978  /
RADECSYS=                'FK5'  /
FILTER  =   'BB#Rc/162_ESO844'  /
AIRMASS =             1.113885  /
EXPTIME =             599.9176  /
EQUINOX =                2000.  /
IMAGEID =                    3  /
GABODSID=                 1401  /
EISID   =                   10  /
OBJECT  =         'BPM16274_3'  /
ZP      =                 -1.0  /
COEFF   =                 -1.0  /
DUMMY0  =                    0  /
DUMMY1  =                    0  /
DUMMY2  =                    0  /
DUMMY3  =                    0  /
DUMMY4  =                    0  /
.
.
END
```

The IMAGEID, GABODSID and EISID are unique identifiers for the chip position within the mosaic, the night of the observation and the image respectively. ZP and COEFF will finally contain magnitude zero point and extinction coefficient. At the end we introduce 20 DUMMY keywords allowing the user to transfer important information from the original headers. For an explanation of the rest of the keywords see Hanisch et al. (2001) and Greisen & Calabretta (2002).