

Intra-Pixel Sensitivity Variation and Charge Transfer Inefficiency — Results of CCD Scans

Hiroyuki Toyozumi^A and Michael C. B. Ashley^{A,B}

^A Department of Astrophysics and Optics, University of New South Wales, Sydney NSW 2052, Australia

^B Corresponding author. Email: m.ashley@unsw.edu.au

Received 2005 March 7, accepted 2005 June 21

Abstract: The efficiency with which a charge-coupled device (CCD) detects photons depends, amongst other factors, on where within a pixel the photon hits. To explore this effect we have made detailed scans across a pixel for a front-illuminated three-phase EEV05-20 CCD using the standard astronomical B, V, R, and I colour filters. Pixel response functions and photometric sensitivity maps are derived from the scan images. Nonlinear charge transfer inefficiency (CTI) effects were observed and corrected for. The resulting images clearly show the intra-pixel sensitivity variations (IPSVs) due to the CCD electrode structure, and its dependence on wavelength. We briefly comment on the implications of IPSVs and CTI for high-precision photometry and astrometry.

Keywords: techniques: photometric — instrumentation: detectors

1 Introduction

In an astronomical context, it is well known that intra-pixel sensitivity variations (IPSVs) in charge-coupled devices (CCDs) can be a significant source of photometric error when the image on the CCD has spatial variations on the sub-pixel level, e.g., if the stars are undersampled. IPSV also leads to errors in astrometry and, in the case of spectra, errors in intensity and centre wavelength assignments of spectral features.

Jorden, Deltorn, & Oates (1994) measured CCDs at multiple wavelengths and plotted the spatial variation pattern and its wavelength dependence. Their results showed that the variation pattern was more complex than would have been naively thought, with significant structure in both the column and row directions. They suggested that the IPSV could be a problem not only for undersampled images but also for spectra.

Kavaldjiev & Ninkov (1997, 1998, 2001) scanned a two-phase front-illuminated CCD, a KAF 4200, with a small light beam, 0.4–0.5 μm in diameter, in two wavelengths (488 nm and 633 nm) allowing them to determine the IPSV to high resolution. They showed that the IPSV pattern of a given pixel was similar to that of neighbouring pixels, although there were differences in detail.

Piterman & Ninkov (2000, 2001, 2002) scanned a thin, back-illuminated CCD, a SITe–502A, using a light beam with 1.7–3.1 μm full-width at half-maximum (FWHM) and B, V, I, and 470 nm narrow-band filters. They showed that the IPSV was smoother, of lower amplitude, and with reduced wavelength dependence when compared with that for front-illuminated devices; they attributed this to the

absence/presence of the gate electrodes. They found a variation in the apparent optical centre of pixels depending on wavelength for both front- and back-illuminated CCDs. They claimed a shift of as much as 30% of the pixel pitch between 905 nm and 400 nm wavelengths for the SITe–502A.

In this paper, we show results from scanning an EEV05-20 CCD using a beam with $<4 \mu\text{m}$ FWHM in four colours (B, V, R, and I). The CCD is front-illuminated and has a three-phase structure with 770×1152 square pixels of side 22.5 μm . The camera is a Wright Instruments model 2 used with the Automated Patrol Telescope (Carter et al. 1992), which is operated by the University of New South Wales at Siding Spring Observatory, Australia. The scans were made with the CCD camera mounted in a test jig in our laboratory.

The CCDs scans were performed by drifting the beam across the CCD and capturing many sub-frame images, i.e., 12×12 pixel regions centred on the beam; the use of sub-frame readout increased the speed of the scanning process and allowed many scans to be averaged in order to increase the signal-to-noise (S/N) ratio.

A complete scan of the CCD consisted of a two-dimensional matrix of the individual sub-frame images (which we call ‘scan frames’) obtained at a regular grid of x, y coordinates. Each scan frame contains a single 12×12 pixel image of the light beam as recorded by the CCD.

From our data, two kinds of images can be extracted: pixel response function (PRF) images and photometric sensitivity map (PSM) images.

A PRF image shows the fraction of the incoming light beam that is detected by a single pixel depending on the relative position of the centre of the beam with respect to the centre of the pixel. An ideal PRF would have a value of 1.0 in a square region of size 1 pixel, and 0.0 elsewhere.

A PSM image shows the total integrated signal recorded by all pixels on the CCD as a function of the x, y coordinate of the light beam with respect to the centre of a reference pixel. In practice, the 12×12 grid of pixels we used captured all but a negligible fraction of the light. An ideal PSM image would be 1.0 everywhere.

The PSM image is convenient for directly determining the photometric error that will result from not correcting for the IPSV effect. If we assume that every pixel is identical, then the PSM image can be derived from the PSF image; it is not possible to go from PSM to PRF.

2 Experimental Setup

There are two main components of our laboratory test setup: a light source with lens, filter, and optical fibre, on a computer-controlled X–Y stage (see Figure 1), and a CCD camera with a lens to produce an image of the light source on the CCD (see Figure 2).

We used a halogen bulb as the light source, and focussed it using an $f/1$ lens placed so as to deliver an approximately $f/3$ beam onto a short length of an optical fibre with a $100 \mu\text{m}$ diameter core, thereby roughly matching the fibre's acceptance angle. A colour filter placed between the lens and the fibre allowed us to control the wavelength range. The halogen bulb was powered by a stable DC power supply to minimise fluctuations in beam intensity during the scanning process. The stability of the light intensity was better than 1% over periods from 1 to 1000 s.

The output end of the optical fibre acts as a $100 \mu\text{m}$ disc of light. We employed an $f/1.1$, 75 mm focal length Rodenstock XR-Heligon lens to produce a de-magnified image of the disc onto the CCD (see Figure 2). By using a 6 m distance between the light source and the lens we achieved a factor of ~ 80 de-magnification, which should give us a $\sim 1.25 \mu\text{m}$ diameter image on the CCD. While the lens is doubtless not perfectly optimised over the full 400–900 nm wavelength range of interest, the use of colour filters made this less critical. There will also be small changes in the illumination of the lens during a scan, leading to changes in the spot shape, but these changes should be negligible over the small area we scanned. In practice, we found that the optical quality of our spot was limited by the field-flattening lens that is permanently mounted as the CCD cryostat window. To reduce the influence of the field-flattener we used a 20 mm diameter aperture mask, as shown in Figure 2, to increase the f number of the beam to 3.75.

The factor of ~ 80 de-magnification of the spot also reduces the precision requirements on the X–Y stage, and we were able to return the beam to within $\sim 1 \mu\text{m}$ of its original position following a scan.

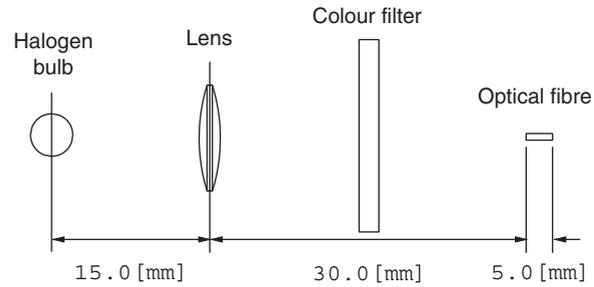


Figure 1 Layout of the light source for the laboratory test setup. All these components are mounted on an X–Y stage.

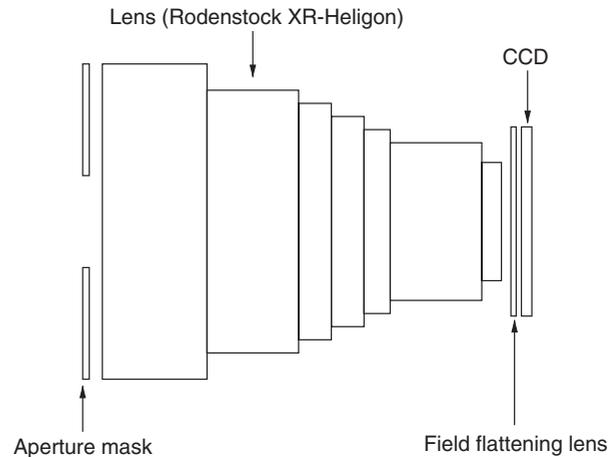


Figure 2 Layout of the CCD camera and lens.

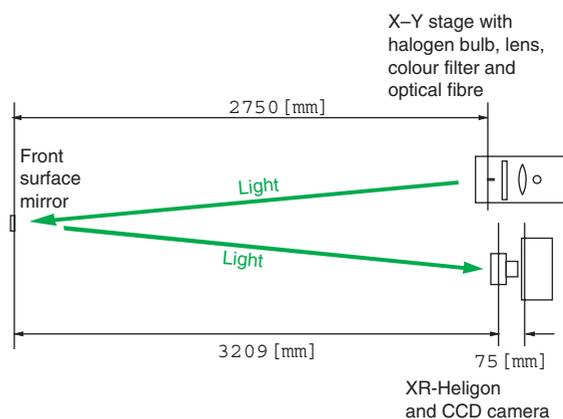


Figure 3 Geometric layout of the test setup.

The shutter of the CCD camera was left open during the scans since the lens protruded into the shutter mechanism. We used electronic shuttering (i.e., moving the electrons comprising the CCD image rapidly away from the light source at the end of each exposure) and chose a low light intensity to ensure that there was negligible smearing of the image as the electrons were being moved.

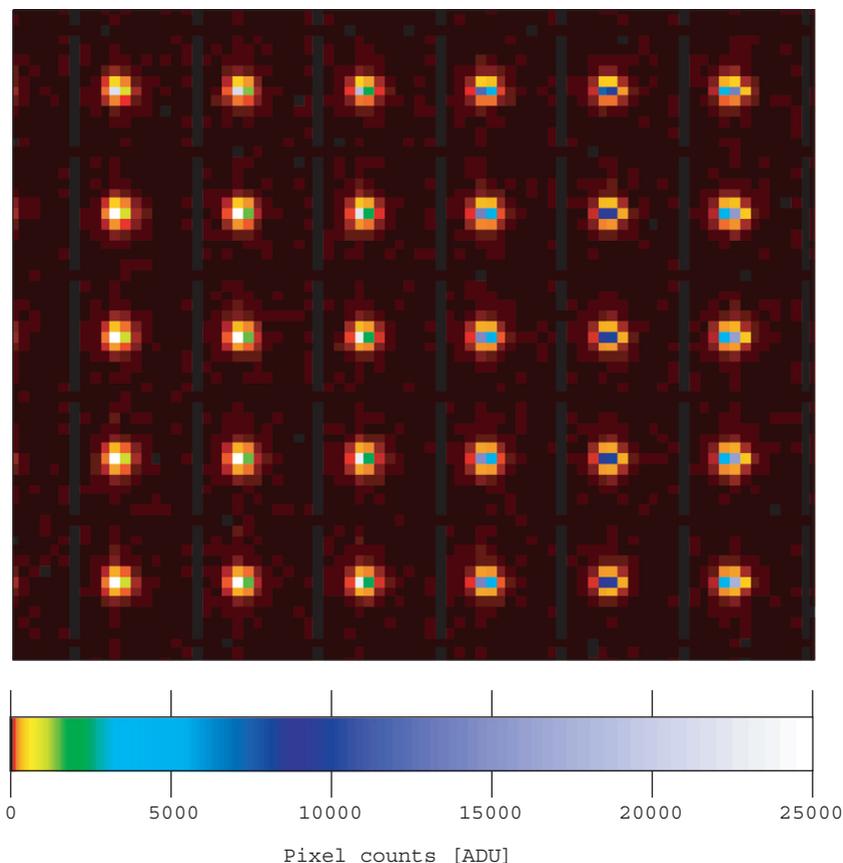


Figure 4 A mosaic of CCD scan frames in V colour. The size of each frame is 12×12 pixels and the stretch of the image is logarithmic.

Figure 3 shows the overall geometric layout of our test setup. We used a front surface mirror to fold the 6 m light path into the available space in our laboratory. The light emitted from the optical fibre was reflected by the mirror, collected by the XR-Heligon lens and focussed on the CCD. The light path was entirely enclosed in blackened tubes, with baffling at various points, to reduce stray light.

3 Scanning the CCD

The CCD was scanned over an approximately 4×4 pixel area in the central region of the CCD. In what follows we refer to the horizontal axis, or rows, of the CCD as the x direction; the rows are parallel to the readout register and perpendicular to the channel stops. Similarly, the vertical axis, or columns, of the CCD are in the y direction.

We scanned the CCD using a raster-like pattern in two different modes: one with x varying most rapidly and the other with y varying most rapidly. A scan line consisted of 25 steps per pixel for a total of 98 steps. Perpendicular to a scan line we made 12.5 steps per pixel for 48 scan lines. At each step on each scan line we exposed for 0.7 s, and captured a 12×12 pixel ‘scan frame’ centred on the scan region. Figure 4 shows a mosaic of individual scan frames for a V band scan after background subtraction.

All scan frames in a raster sequence were stored in a data file on a computer. To assist with later analysis, exposure information was inserted into the data by overwriting the lower-left four pixels (8 bytes) of each scan frame. The information consisted of a continuous frame number (2 bytes), an exposure start time stamp in microseconds (3 bytes), and an exposure end time stamp (3 bytes). The continuous frame number and time stamps were used to check for any loss of synchronisation between the computer programs controlling the X–Y stage control and the CCD.

The intensity of the beam for each scan was carefully adjusted to keep a high S/N ratio while avoiding saturation and nonlinearity. The maximum pixel count measured during the scans was $\sim 20\,000$ ADU. The gain of the CCD camera was 10 electrons per ADU, so the S/N ratio from Poisson statistics was ~ 450 at the peak.

The CCD operating temperature was regulated at $\sim -40^\circ\text{C}$ to reduce dark current (< 0.8 electron per pixel per second).

We could not measure the exact size of the beam spot directly. One method we considered was to magnify the beam using an additional lens and then observe the magnified spot with the CCD. However, because of the field-flattening lens, the original beam profiles projected on the CCD surface were not simply the scaled images

of the magnified profiles. Instead, we estimate the beam profiles by fitting a theoretical model to the derived PRF images.

CCD scans were made in four colours: B, V, R, and I, defined by filters made according to the Bessell formulation (Bessell 1990). The light intensity was varied when swapping filters to keep the signal detected by the CCD approximately the same. Table 1 shows the colours, the approximate central wavelength, approximate bandpass, and the number of scans performed.

4 Data Reduction

The procedure for data reduction and calibration was as follows:

1. subtract a constant DC bias from all pixels,
2. apply a linearity correction derived from the known performance of the CCD readout amplifier,

Table 1. Filters and the number of scans

Colour	Central wavelength (nm)	Bandpass (nm)	Scans	
			horizontal	vertical
B	440	100	4	5
V	550	90	4	5
R	650	150	4	4
I	800	140	3	3

3. normalise the intensity of each frame by correcting for small variations in the exposure time (as recorded in the individual scan frames), and
4. fit and subtract background values.

Inter-pixel sensitivity calibration (i.e., flat field correction) was omitted since we were interested in measuring differences from one pixel to the next.

4.1 Deriving the PRF and PSM Images

Since our scans of the CCD covered 3×3 pixels completely, we could obtain nine PRF images from each scan. Furthermore, since the CCD has been scanned in two directions with different scanning resolutions there were two sets of PRF images: x PRF images with 98×48 pixels and y PRF images with 48×98 pixels, see Figure 5.

To maximise the fidelity of our derived PRF we took considerable care in combining the nine sets of PRFs in x and y . The two orthogonal sets of scans also assisted in correcting for fluctuations in the light intensity. We constructed PRFs that best represented all the data and achieved the full $1/25$ pixel resolution in both directions, see, e.g., Figure 6. The final PRFs, after correction for a ‘ghost’ effect discussed in the next section, are shown in Figures 11–13.

The PSM images were derived in a similar fashion and are shown in Figures 14 and 15.

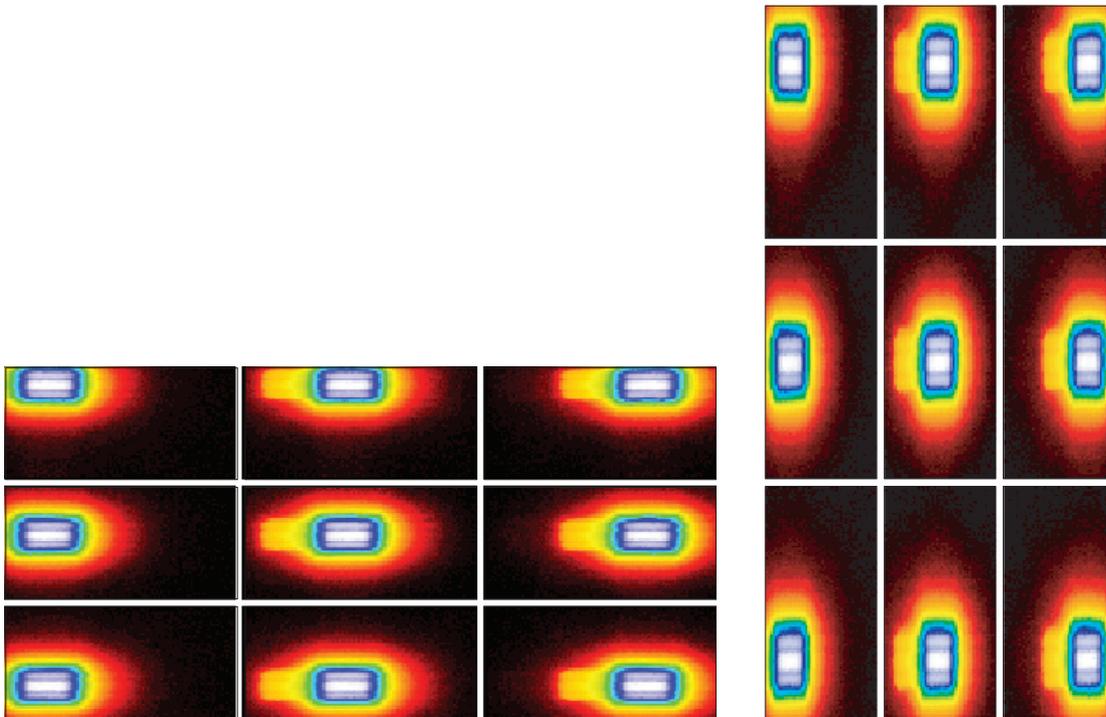


Figure 5 The nine independent PRF images from an x scan (left) and a y scan (right). The stretch of the images is logarithmic and each PRF covers 4×4 pixels.

5 Dealing with Ghosts

5.1 The Mystery of the Ghosts

The PRF images in Figure 5 show an anomalous extension, or ‘ghost’, to the left of each central peak. This effect is more clearly seen in the combined PRF (see Figure 6).

Since the ghost appears to be a duplicate of the primary pattern, placed one pixel to the left, we suspected that it was due to an electronic phenomenon rather than an optical one. The most likely cause is charge transfer inefficiency (CTI) in the electron transfer in the readout register. When charge stored in a pixel in the readout register is shifted by one pixel towards the readout amplifier, a small amount of charge is left behind (Janesick 2001, p. 387). The observed intensity of our ghost image is consistent with a readout register CTI of 5×10^{-5} , which is a realistic value for this CCD.

Note that in our scan images, a ghost to the left corresponds to charge trailing to the right in the readout register, which is consistent with CTI being the cause.

A second possibility is dielectric absorption in the readout electronics of the CCD camera. Dielectric absorption is a voltage re-bounce phenomenon associated with capacitors when they are completely discharged after having been charged for a relatively long time. In the case of CCD camera systems, charges transferred from the CCD are integrated in a capacitor as part of the analog to digital conversion process, and dielectric absorption can bias the counts recorded for downstream pixel values.

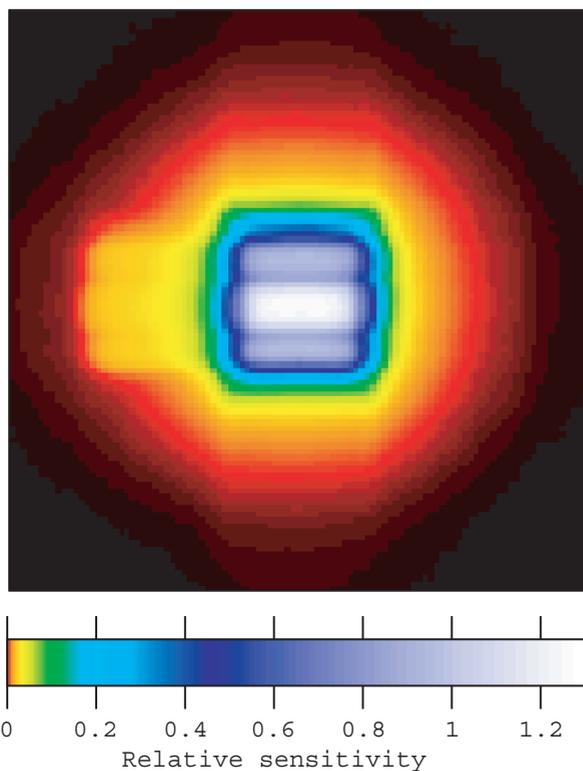


Figure 6 A ghost (the apparent duplicate of the main pattern, one pixel to the right of the centre) appearing in the B band scan image. The stretch of the image is logarithmic and the size corresponds to 4×4 pixels.

Another possible cause is a remnant signal due to the slow discharging of a post-amplifier filter in the readout circuit. In a CCD camera system, this may happen when pixel values are read rapidly compared to the time constant of a post-amplifier filter; the effect is most noticeable when dark pixels follow a very bright pixel (Janesick 2001, p. 582).

Later investigations showed that the magnitude of the ghost effect increased linearly with the x coordinate of the image. Of the three possible causes discussed above, only CTI is consistent with this observation.

5.2 Removing the Ghosts

We arrived at a simple empirical method for measuring and removing the CTI effect by thinking about what happens as a well-focussed light beam traverses pairs of pixels horizontally from left to right (see Figure 7). First, the value of the left pixel increases and reaches a maximum value, leaving the right pixel value close to zero; then, the value of the left pixel decreases and approaches zero while the value of the right pixel increases and approaches the maximum value; finally, the value of the right pixel decreases leaving the left pixel value zero. If the observed signals from the left and right pixels are plotted against each other, as in the top-left frame of Figure 8, they produce a pattern which depends on the PRF of the CCD and the beam profile.

If the light beam is perfectly symmetrical from left to right, the pixel pair plot should be symmetrical about a 45° line, and this is close to being true in Figure 8. However, if you examine the regions close to the axes (enlarged in the top-right and bottom-left frames in the figure), you see an obvious asymmetry. If we make the assumption that the bottom-left frame represents the true shape of the PRF edge, then the difference between this and the top-right frame can be attributed to CTI.

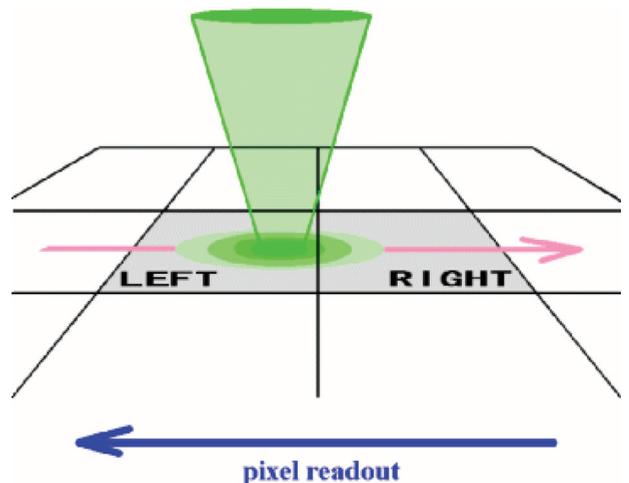


Figure 7 A pixel pair and an incoming light beam. Each square grid represents a pixel. The arrow shows the line along which the beam traverses.

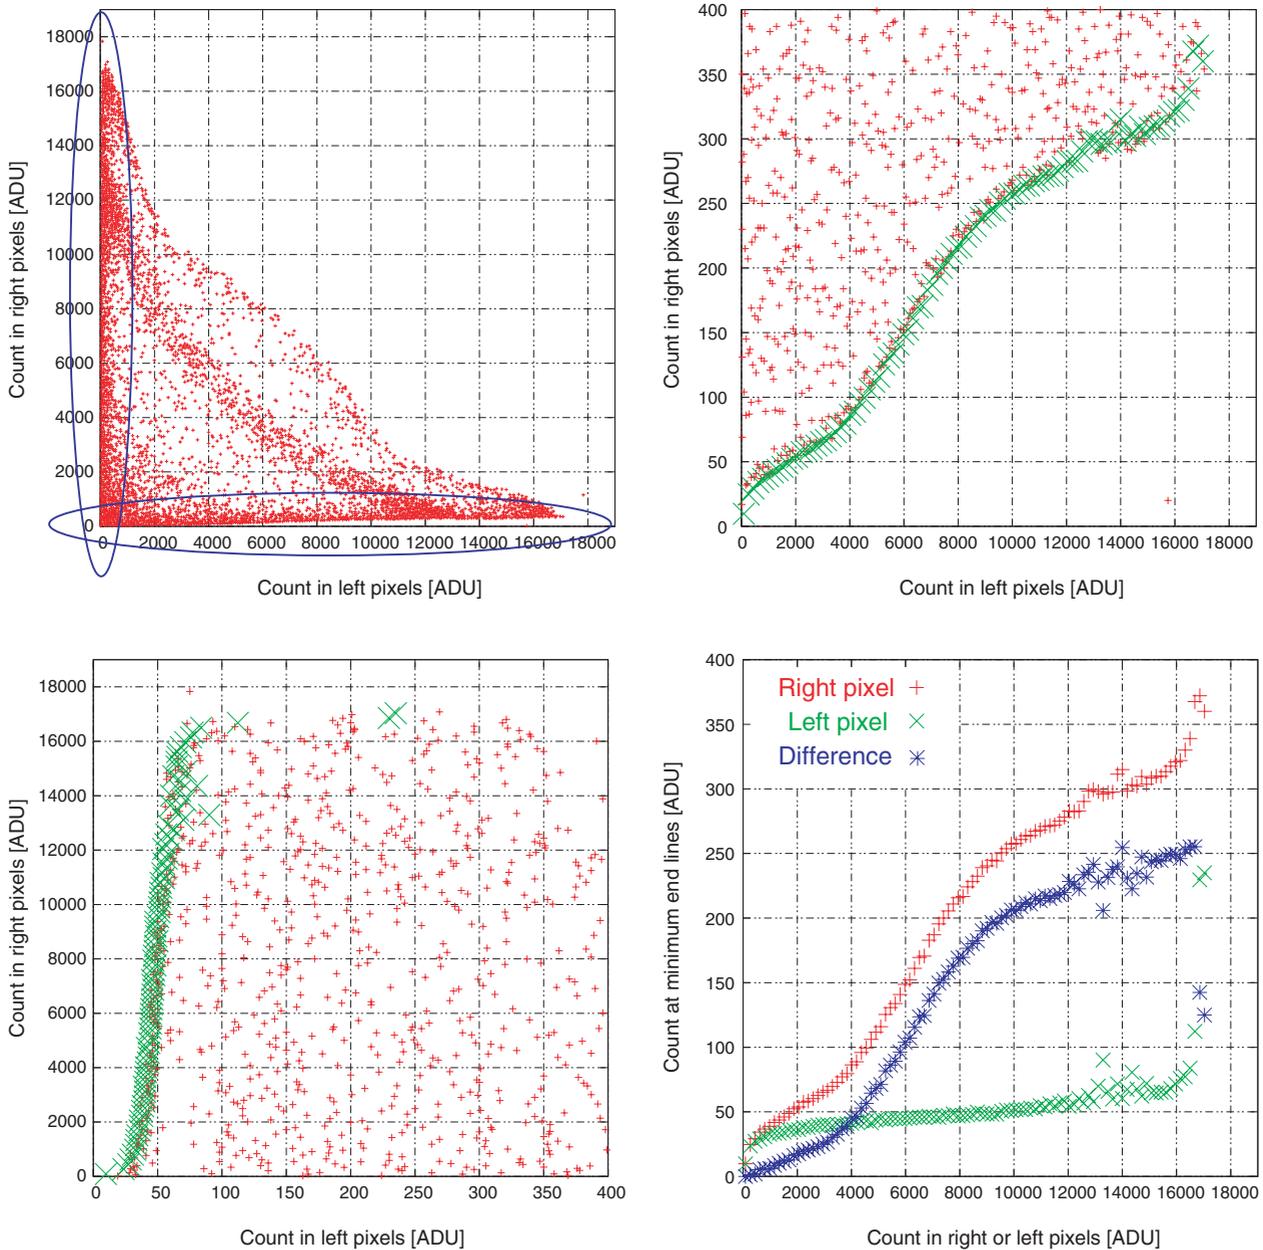


Figure 8 Extraction of the ghost effect. The left-top figure shows a plot of all pixel pairs in a scan for the B band. The right-top and left-bottom figures are magnified plots of the minimum end lines circled in the left-top figure. The points presented with (\times) are the averaged points that were calculated by gathering points near the minimum end lines. The right-bottom figure shows a re-plot of the averaged points.

The bottom-right frame in Figure 8 shows a re-plot of the relevant regions in the pixel-pair plot, showing only the pixels bounding the regions. The difference in the two boundary lines is the CTI effect.

To gain further insights we repeated the pixel-pair plotting for all four observed wavebands. The results are shown in Figure 9. The close agreement between the wavebands is expected from the CTI explanation for the ghosts and argues against a contribution from an asymmetry in the illuminating beam or the CCD pixel response function itself, both of which would be expected to show colour dependence.

The nonlinear nature of the CTI correction curve has been commented on before (Banghart et al. 1991; Marshall & Marshall 2003), and the trend we see for the curve to flatten at higher signal levels is expected. To our knowledge, however, the S-shaped curve at lower levels has not been reported before.

Interestingly, we have not seen any evidence of CTI effects in the vertical clocking, despite the fact that parallel CTI is usually worse than that for the readout register.

CTI effects are not obvious in typical astronomical images since the instrumental point spread function (PSF)

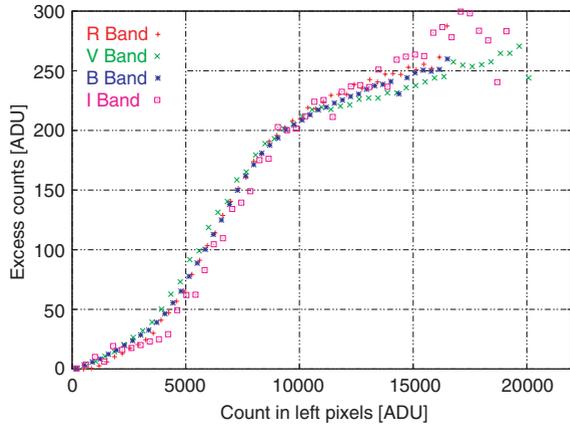


Figure 9 Excess counts caused to downstream pixels due to charge transfer inefficiency for all four observed wavebands.

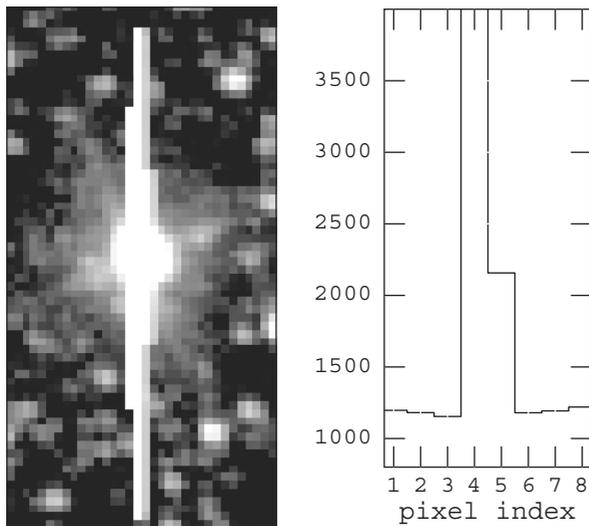


Figure 10 The CTI effect seen in an astronomical image of a saturated star. Pixels to the right of the saturated column have increased counts due to the effect of charge transfer inefficiency in the read-out register. The right-hand figure shows a horizontal cut through the image, 9 pixels down from the top. The saturated pixel in this plot, at an x coordinate of 529, had 40 000 counts, 1000 of which were redistributed by CTI to the pixel to the right. This implies a horizontal CTI of 5×10^{-5} .

is usually not sharp enough to allow the effect to be easily seen. However, the effect can be readily detected in pixels downstream of the column overflow from a saturated star. Figure 10 shows this effect for our CCD.

We used the curve in Figure 9 to subtract the CTI effect from the raw scan frames, and achieved a satisfactory result as can be seen in Figures 11–13.

6 Results

Figure 11 shows the final PRF images, Figure 12 shows three-dimensional plots of the PRF images, and Figure 13 shows cross-sections of the PRF images through the centre of a pixel. Figure 14 shows the final PSM images, and Figure 15 shows cross-sections of the PSM images.

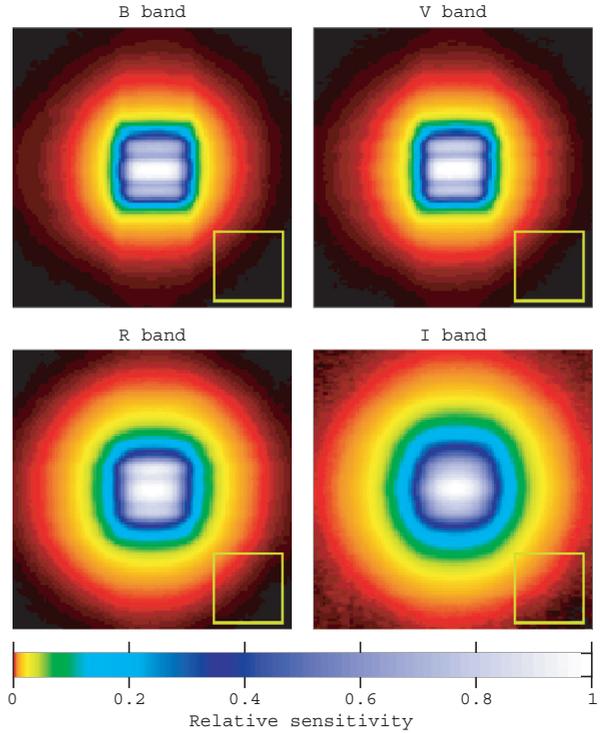


Figure 11 The two-dimensional plot of the pixel response functions after correction for charge transfer inefficiency. The stretch of the images is set logarithmic and the size corresponds to 4×4 pixels. The white rectangles in the images show the size of a CCD pixel.

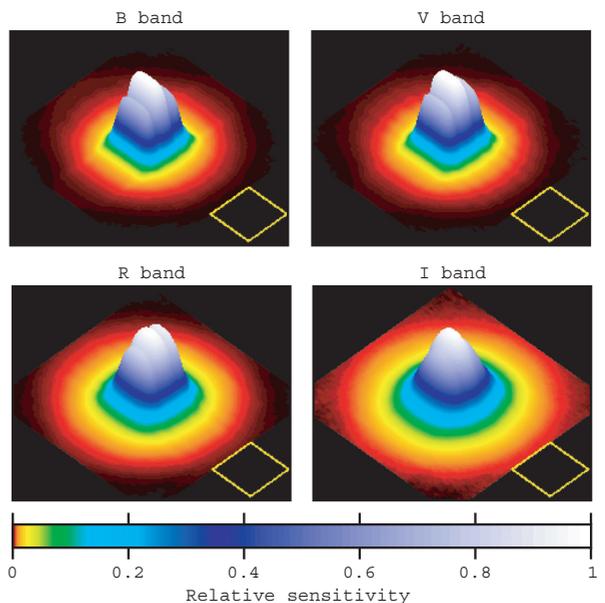


Figure 12 The three-dimensional plots of the pixel response functions after correction for charge transfer inefficiency. The height of the surface is linear with sensitivity. The size corresponds to 4×4 pixels. The white rectangles in the images show the size of a CCD pixel.

We can estimate an upper limit on the FWHM of our light beam by assuming it is Gaussian in profile, and then examining the maximum spatial derivative in the observed PRFs. Table 2 shows these limits. If the

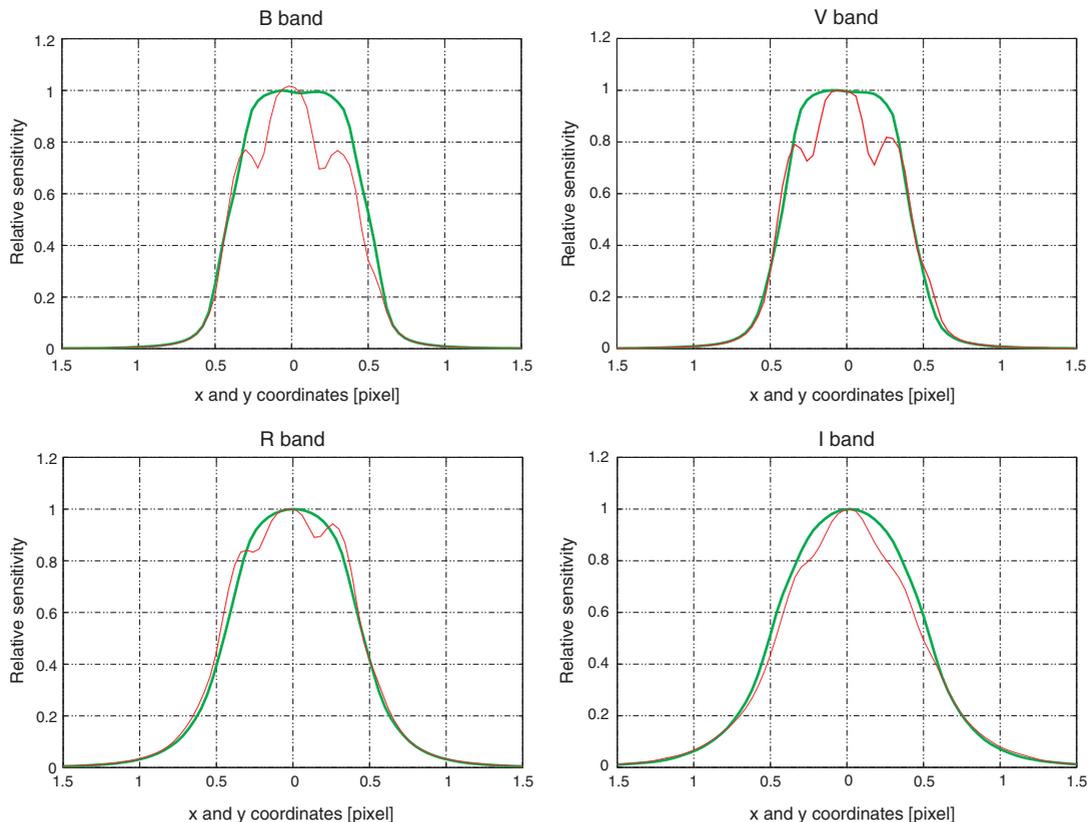


Figure 13 Cross-sections of the pixel response functions. The thick lines show the cross-section in the x direction and the thin lines show the cross-section in the y direction. The cross-sections were chosen to pass through the peak signal in the pixel.

beam was larger than this in size, then the observed PRFs would be smoother. Note that the PRF is expected to become smoother at longer wavelengths since the photons penetrate deeper into the silicon, and consequently have a greater chance of diffusing into neighbouring pixels; therefore the I band beam diameter in Table 2 is a very conservative upper limit.

7 Interpretation

A three-phase CCD, such as the EEV05-20 studied here, uses three horizontal electrodes per pixel. The electrodes are deposited on the surface of the CCD and hence directly attenuate any light which falls on them. This effect is clearly seen as regions of reduced sensitivity in the PRF images in Figure 11.

In the y direction, the channel stops define the extent of the charge that is collected under each pixel, and there is no evidence for any other structure in y in our PRF images.

The broad low-level wings that we see in the PRF could partly be due to the profile of our light beam, although the fact that the wings become broader as the wavelength is increased is an expected consequence of the increased depth at which the longer wavelength photons are absorbed.

The PSM images in Figures 14 and 15 show sensitivity peaks on channel stops as previously reported by Jorden et al. (1994). Jorden et al. attributed this to the

charge diffusion due to the deep penetration depth of I band photons. The PSM images show no evidence for significant differences in the sensitivity pattern from one pixel to another.

8 The Effects on Astronomical Images

To see the effect of non-ideal PRFs on an astronomical image, it is necessary to convolve the laboratory PRF with the instrumental point spread function (iPSF) of the telescope. More subtly, it is necessary for the laboratory measurements of the PRF to have been made using a light beam with the same f/l number as the telescope. This latter point is important: IPSV depends on the f/l number since the exact details of the reflection and refraction in the surface layers of the CCD depends on the incident angle of the photons. The test beam should mimic the telescope closely, even including central obstructions and vignetting, since these alter the distribution of angles of the photons. There are also likely to be changes in the angular distribution of the photons in the image plane of the telescope, so the effect of IPSV can vary across the CCD, leading to systematic errors in photometry and astrometry. These points have not been widely appreciated in the literature.

Table 3 shows the photometric errors in magnitudes that would occur if our CCD was used in a telescope that could produce star images as well focussed as our laboratory

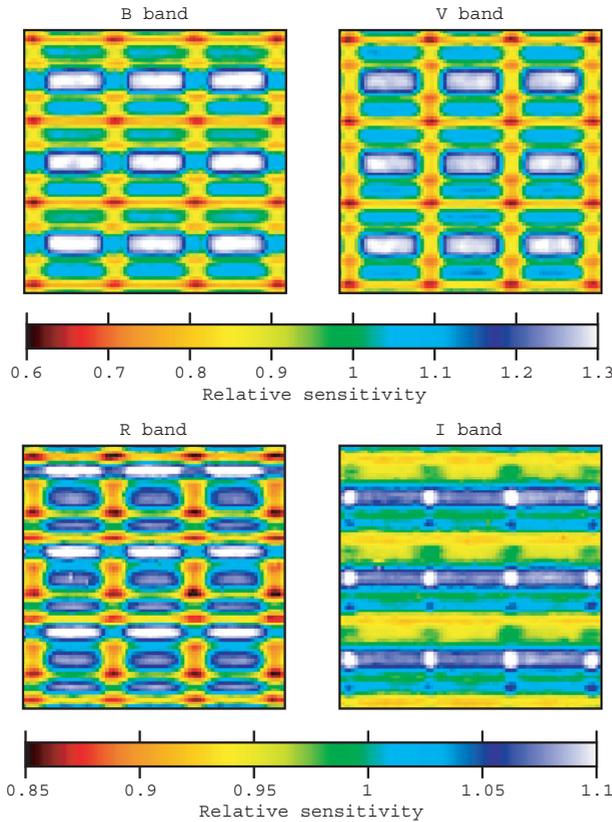


Figure 14 The two-dimensional plots of the photometric sensitivity maps after correction for charge transfer inefficiency. The stretch of the images is set linear, and the image size corresponds to 3.2×3.2 pixels.

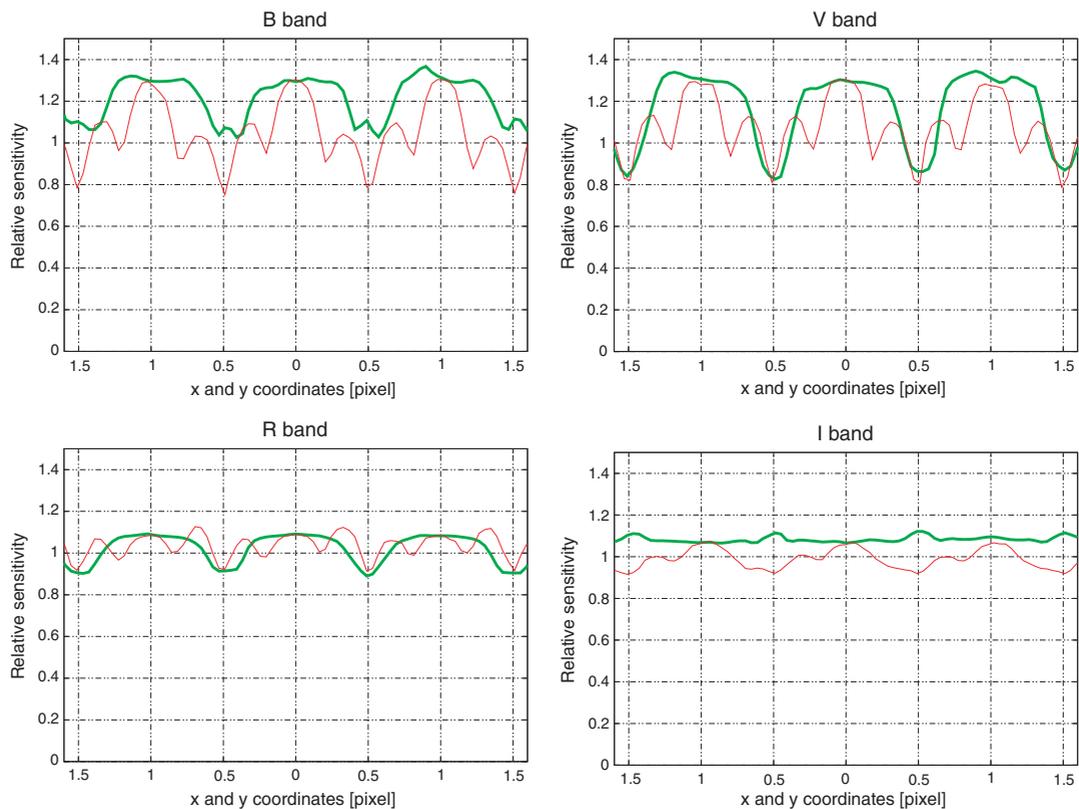


Figure 15 Cross-sections of the photometric sensitivity maps. The cross-sections were chosen to pass through the centre of the central pixels in Figure 14. The thick lines show the cross-section for the x direction and the thin lines show the cross-section for the y direction.

setup. In reality, star images are affected by diffraction and seeing, leading to a typical FWHM of between one and two pixels. While this reduces the photometric errors significantly, they can still easily reach the 10 millimag level, which is significant for many high-precision astronomical programmes.

To the first order, charge transfer inefficiency can be considered as a small redistribution of electrons to neighbouring pixels, and so should have little effect on aperture

Table 2. Upper limits on the scan beam diameter

Colour	FWHM (pixels)
B	0.15 (3.4 μm)
V	0.15 (3.4 μm)
R	0.15 (3.4 μm)
I	0.34 (7.6 μm)

Table 3. Photometric errors resulting from the pixel response function

Colour	Magnitude uncertainty
B	± 0.30
V	± 0.28
R	± 0.12
I	± 0.10

photometry even if CTI is nonlinear. However, it will alter the apparent point spread function of star images depending on their x , y coordinates and brightness, and so will have small second-order effects on photometry and astrometry. Furthermore, if CTI is not corrected for prior to flat-fielding, then the redistributed electrons will be multiplied by the sensitivity of the pixel that they are in, rather than the pixel from where they came. And to further complicate the situation, the flat field itself would have been affected by CTI. Both effects are likely to be negligible provided that the flat field is fairly flat. In any case, it is straightforward to make an adequate correction to the raw images if the CTI curve (Figure 9 in our case) is known.

9 Conclusion

With a relatively simple laboratory setup we have readily made measurements of the pixel response function of a CCD using standard astronomical filters. We have also measured the nonlinear charge transfer inefficiency of the CCD.

The pixel response function of a pixel is not purely a function of the CCD and the wavelength, but depends on the angular distribution of the incoming photons (i.e., the $f/\#$ number and any effects due to obstructions and

vignetting). This can lead to subtle systematic effects in photometry and astrometry depending on the x , y position on the CCD.

Acknowledgments

We greatly appreciate advice from a number of people on the CCD-world mailing list. Jon Everett assisted us with the optical test setup in our laboratory.

References

- Banghart, E. K., Lavine, J. P., Trabka, E. A., Nelson, E. T., & Burkey, B. C. 1991, *ITED*, 38(5), 1162
- Bessell, M. S. 1990, *PASP*, 102, 1181
- Carter, B. D., Ashley, M. C. B., Sun, Y.-S., & Storey, J. W. V. 1992, *PASA*, 10, 74
- Janesick, J. R. 2001, *Scientific Charge-Coupled Devices* (Bellingham, WA: SPIE Press)
- Jorden, P. R., Deltorn, J.-M., & Oates, A. P. 1994, *Proc. SPIE*, 2198, 836
- Kavaldjiev, D., & Ninkov, Z. 1997, *Proc. SPIE*, 3019, 42
- Kavaldjiev, D., & Ninkov, Z. 1998, *OptEn*, 37, 948
- Kavaldjiev, D., & Ninkov, Z. 2001, *OptEn*, 40, 162
- Marshall, C. J., & Marshall, P. W. 2003, http://radhome.gsfc.nasa.gov/radhome/papers/CCD_Lessons_Learned.pdf
- Piterman, A., & Ninkov, Z. 2000, *Proc. SPIE*, 3965, 289
- Piterman, A., & Ninkov, Z. 2001, *Proc. SPIE*, 4306, 280
- Piterman, A., & Ninkov, Z. 2002, *OptEn*, 41, 1192