CHAPTER 4: CCD Performance

In This Chapter...

4.1 Introduction

The WFPC2 CCDs are thick, front-side illuminated devices, with a format of 800x800, 15x15µm multi-pinned phase (MPP). MPP allows CCD exposure with the total inversion of all phases. The Si-SiO_2 interface, at the surface of the CCD, is pinned at the substrate potential, directing signal charge away from the Si-SiO_2 interface states towards the buried n-channel. Figure 4.1 shows a schematic which illustrates the principle of MPP (modified from Janesick et al. 1989). The front-side Si-SiO_2 interface significantly affects the performance of CCDs, so MPP operation yields many practical benefits including reduced dark noise, better charge transfer efficiency (CTE), rapid removal of residual images, excellent pixel-to-pixel uniformity, and improved radiation hardness. MPP technology has been demonstrated and characterized in both Loral (Janesick, et al., 1989) and Tektronix devices (Woodgate, et al., 1989). The CCD sensors for WFPC2 were made by Loral in 1991 and processed and packaged for flight at JPL.

The Loral CCDs are illuminated from the 'front' surface, i.e., the light passes through the polysilicon gate structure overlying the 10µm thick active silicon layer. Because the WFPC2 devices are front-side illuminated and supported by a bulk silicon substrate, the CCD surface is flat, which has reduced the uncertainties in the astrometric calibration to about the 1/10 pixel level.

Figure 4.1: MPP Operating Principle. A schematic showing the ideal potential profile through a frontside illuminated CCD whose front surface is inverted with multi-pinned phase (MPP). The CCD consists of a polysilicon gate, which forms part of the electrode structure, a surface layer of oxidized silicon (SiO_2) and the epitaxial layer which comprises p-doped silicon with an n-doped buried-channel for charge transfer. MPP pins the surface potential by populating the Si-SiO_2 interface with holes. The holes passivate the Si-SiO_2 interface states and create an electric field which directs signal charge away from the interface towards the buried n-channel.

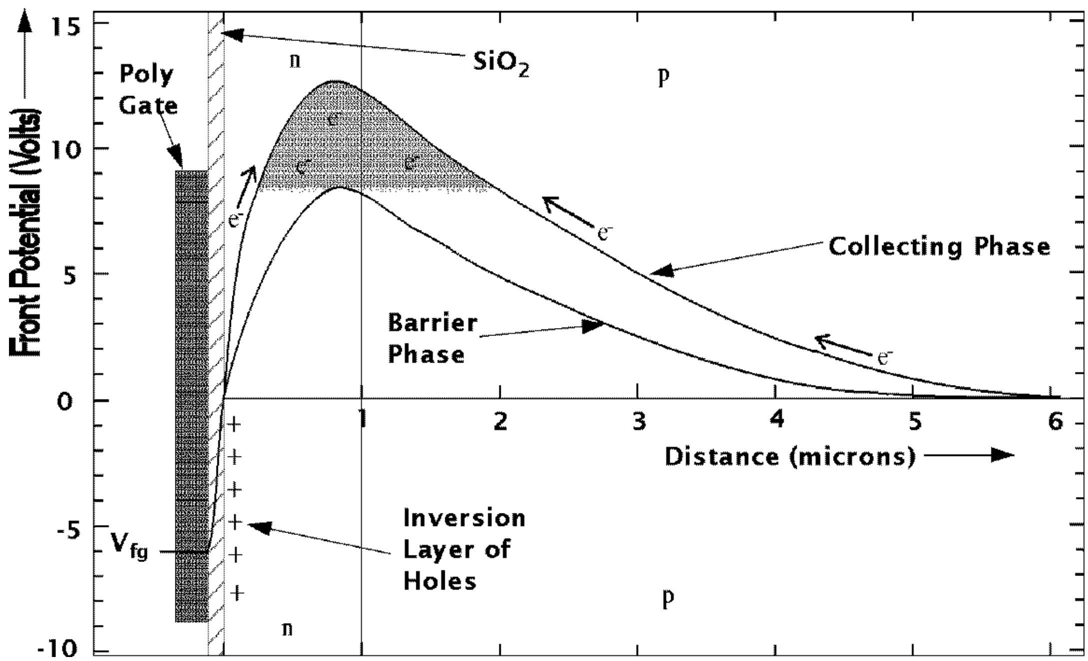

In this section the performance of the WFPC2 CCDs is reviewed, and compared to the WF/PC-1 devices. A summary of device characteristics is given in Table 4.1.

4.2 Quantum Efficiency

The Loral CCDs are thick, front-side illuminated devices. This lowers their intrinsic QE, due to the absorption of incident light by the polysilicon electrode structure on the front-side surface of the CCD. We also note that due to its MPP operation, the QE of the Loral devices is stable without maintenance such as UV flooding.

The front surfaces of the CCDs are overcoated with a Lumogen phosphor, which serves as the primary detection medium for photons shortward of about 4800Å, down-converting these to 5100Å - 5800Å. Its long wavelength cutoff (4800Å) is also well matched to a CCD's intrinsic sensitivity. The QE of the four flight WFPC2 CCDs is shown in Figure 4.2, which demonstrates the uniform UV

Table 4.1: Comparison of WF/PC-1 and WFPC2 CCDs.

Parameter	WF/PC-1[a]	WFPC2
Device	TI	Loral
Architecture	Thinned	Thick
Illumination	back-side	front-side
Format	800×800	800×800
Pixel size	15^2 µm	15^2 µm
UV Phosphor	Coronene	Lumogen
Dark rate	0.03 e^- $pixel^{-1}$ s^{-1}(-87°C)	~0.0045 e^- $pixel^{-1}$ s^{-1}(-88°C)
Read noise	13e^- RMS	5e^- RMS
Full well depth	40,000 e^-	~90,000 e^-
Gain	8e^- DN^{-1}	7e^- DN^{-1} or 14e^- DN^{-1}
ADC range	12 bits (4096 DN)	12 bits (4096 DN)
Full range (e^-)	~30,000e^-	~53,000e^-
QE 6000Å	50%	35%
QE 2500Å	12%	15%
WFC resolution	0.10″ $pixel^{-1}$	0.0996″ $pixel^{-1}$
PC resolution	0.043″ $pixel^{-1}$	0.0455″ $pixel^{-1}$

a. WF/PC-1 data are available through the STScI data archive.

response of 10-15% and a peak optical QE of 40%. This phosphor coating also produces an enhancement of DQE at visual wavelengths, since it acts as an anti-reflection coating

4.3 Dynamic Range

Linear full well capacity for these devices, clocked appropriately for the MPP mode, is approximately 90,000e^- $pixel^{-1}$. Flight qualified ADCs with higher dynamic range (>12 bits) were not available, so WFPC2 operates the two available ADCs at different gain factors, to take partial advantage of both the low read noise and large available full well depth.

One channel has a gain of 14e^- DN^{-1}, which significantly undersamples the CCD read noise (5 e^- $pixel^{-1}$ RMS), and gives a digital full well of about 53,000e^-. The other channel has a gain of 7e^- DN^{-1} which is comparable to the CCD read noise, and saturates at about 27,000e^-. The choice of gain factor will be determined by the scientific objective. The 7 e^- DN^{-1} channel is best suited for faint object and UV imaging, where the lower CCD read noise will be most

Figure 4.2: Pre-flight DQE Measurements on WFPC2 CCDs. The differences between the chips are probably due to systematic measurement error, and do not reflect a real difference in sensitivity.

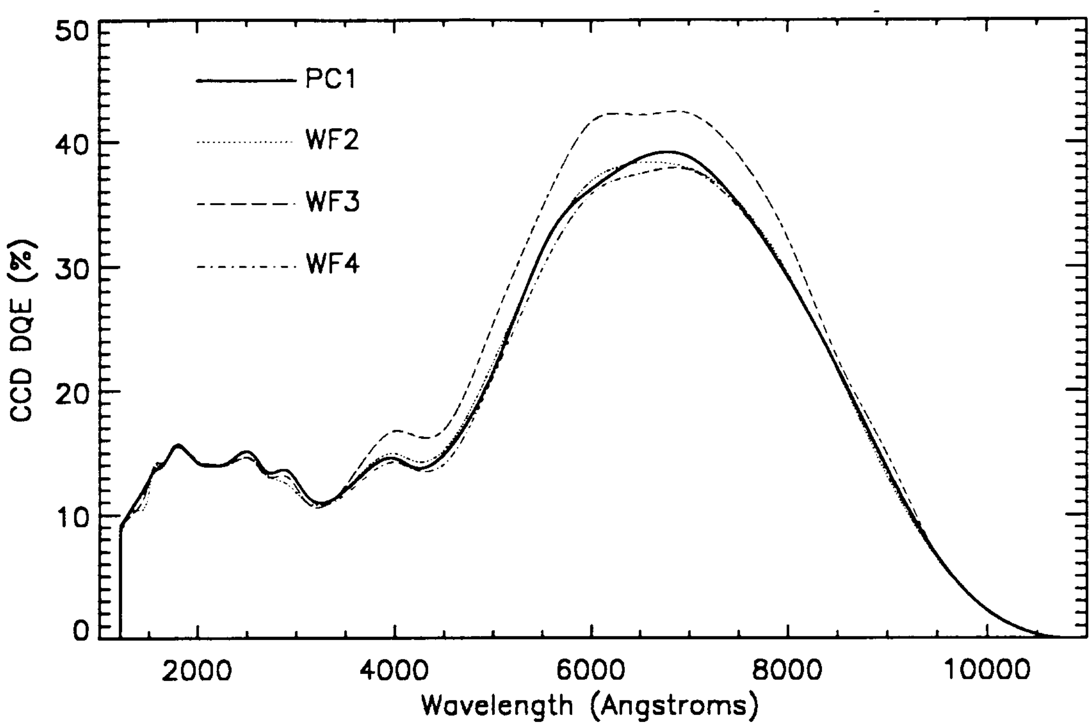

effective. For example, it should be used for UV imaging of planets or narrowband imaging of high redshift galaxies. The 14 e^- DN^{-1} channel has slightly higher effective read noise due to the quantization granularity, but can be used for programs where a signal level in excess of 27,000e^- is required. Even when imaging faint sources, it may be desirable to retain the high signal-to-noise information on brighter field stars as a PSF reference.

Use of the 14 e^- DN^{-1} channel also allows reasonable recovery of counts for isolated, saturated point sources by summing over the saturated pixels (assuming that the charge bleeding does not extend to the edges of the CCD). See Gilliland (1994).

4.4 Bright Object Artifacts

4.4.1 Blooming

Blooming up and down a CCD column occurs when more than about 90,000e^- (the full well capacity) are collected in any pixel. When the pixel is full, the charge will flow into the next pixels along the column, and so on. The orientation of the bloomed column(s) on the sky depends on the readout direction of the particular CCD (see Figure 1.1 on page 2 or Figure 3.11 on page 62) and the roll angle of the spacecraft. This effect is visible in Figure 4.3 which shows a logarithmic stretch of the image resulting from a 100s exposure on a star of V magnitude 2.6 through filter F502N in the PC.

Extreme overexposure of the Loral CCDs is not believed to cause any permanent effects, and therefore the WFPC2 does not have a bright object limit.

The WFPC2 CCDs can be operated in a non-standard mode during the integration phase of an exposure, in order to limit the blooming to only those columns containing the bright sources. This is accomplished by operating the serial transfer register clocks during the integration (using the optional parameter CLOCKS as specified in the Proposal Instructions). See section 2.6, “Serial Clocks”, on page 27 for details.

4.4.2 Horizontal Smearing

During readout of a badly overexposed image, there is spurious charge detected by the readout electronics. The apparent brightness of the stellar halo is higher to the right of the saturated columns. This is particularly obvious at the bottom of the image in Figure 4.3 which is a region in the shadow of the pyramid edge.

The horizontal “smearing” seen in highly saturated images can be modeled as an exponential function which decays over a few rows after a saturated pixel is encountered. The effect itself temporarily saturates after about ten saturated pixels (subsequent saturated pixels have no effect). The effect is twice as bad with gain 7 e^{-} DN^{-1} than with gain 14 e^{-} DN^{-1}. This model only works on very highly saturated stellar images.

In Figure 4.3, the image to the right side of the saturated columns is brighter than the left side; and the brightness increases as the number of saturated columns increases. This effect appears to be a signal which starts at a saturated pixel and decays over the next few rows, wrapping around as it does so. The signal is additive with each successive saturated pixel. Jumps are obvious when the number of saturated columns changes. The problem is a known characteristic of the amplifier electronics, and effort was made to minimize it during design. The increase in signal in rows with saturated pixels is also seen in the over-scan region (the over-scans are provided in “`.x0d`” files from the pipeline)

An exponential function fits the effect reasonably well. An appropriate algorithm creates an array to contain the signal model. It searches through the uncalibrated image (with the over-scan region included) in the sequence in which the pixels are read out. When it encounters a saturated pixel, it adds an exponential function to the model array, beginning at that pixel. The function has the form $s(x)=Ae^{-x/h}$, where x is the offset from the saturated pixel and only positive x values are included. The half-width, h, and amplitude, A, appear to vary from frame to frame and must be determined on the image itself. As more saturated columns are encountered in a row, the signal intensity builds up in the model image. The image can then be “improved” by subtracting the model from the raw image.

The amplitude and half-width parameters can be obtained by trial and error. The typical parameters vary slightly for each chip. The amplitude per saturated pixel is typically 1.75 DN (gain 7) or 0.2 DN (gain 14). On the other hand the half-width at a gain of 14 is larger (h=1800) than at 7 (h=350). So the total integrated effect is about twice as bad at gain 7. A straightforward application of the above algorithm cleaned up most of the signal in rows which had a few saturated columns, but over-subtracted in rows with a large number. The algorithm can be modified to saturate by making the parameter A, which gives the peak contribution from a single saturated pixel, depend on the current level of the effect: $A=A_0*(1-C/C_{max})$. This implies that the correction is never larger than C_{max} no matter how many saturated pixels are encountered. C_{max} is approximately 14 DN for a gain of 7 and 10 DN for a gain of 14.

The algorithm gives improvement only on highly saturated stellar images (where the star is saturated to 3 or 4 columns at the edges of the chip). On less saturated data, it over-subtracts significantly. This indicates that the problem is nonlinear, and therefore a general algorithm applicable to all data will be difficult to develop.

4.4.3 Diffraction Effects and Ghost Images

Several other artifacts that are common in saturated stellar images are also obvious in Figure 4.3. The spider diffraction spikes caused by both the OTA spiders and internal WFPC2 spiders are at 45° to the CCD columns in all cameras.

The halo around the stellar image is well above the diffraction limit in intensity. Also there are ghost images which result from internal reflections in the filters and in the field-flatteners. These topics are discussed fully in the next Chapter.

4.5 Residual Image

Residual images are seen in front-side-illuminated CCDs, and are associated with the front-side $Si\text{-}SiO_2$ surface interface. When the full well is exceeded electrons can become trapped at the $Si\text{-}SiO_2$ interface. This trapped charge is slowly released giving rise to residual images. Inverted phase operation (MPP) allows holes to recombine with the trapped electrons at the front-side interface, and so residual images dissipate in a matter of minutes.

A second potential source of residual images, which occurs only in front-side-illuminated CCDs, is known as residual bulk image (RBI). Long wavelength photons can penetrate deeply enough to produce charge in the substrate. Most of this charge recombines rapidly (due to short carrier lifetimes), but some may diffuse into the epitaxial layer, where it can become trapped in epitaxial interface states. Residual images can occur as this charge is slowly

Figure 4.3: Saturated Stellar Image Showing Horizontal Smearing.

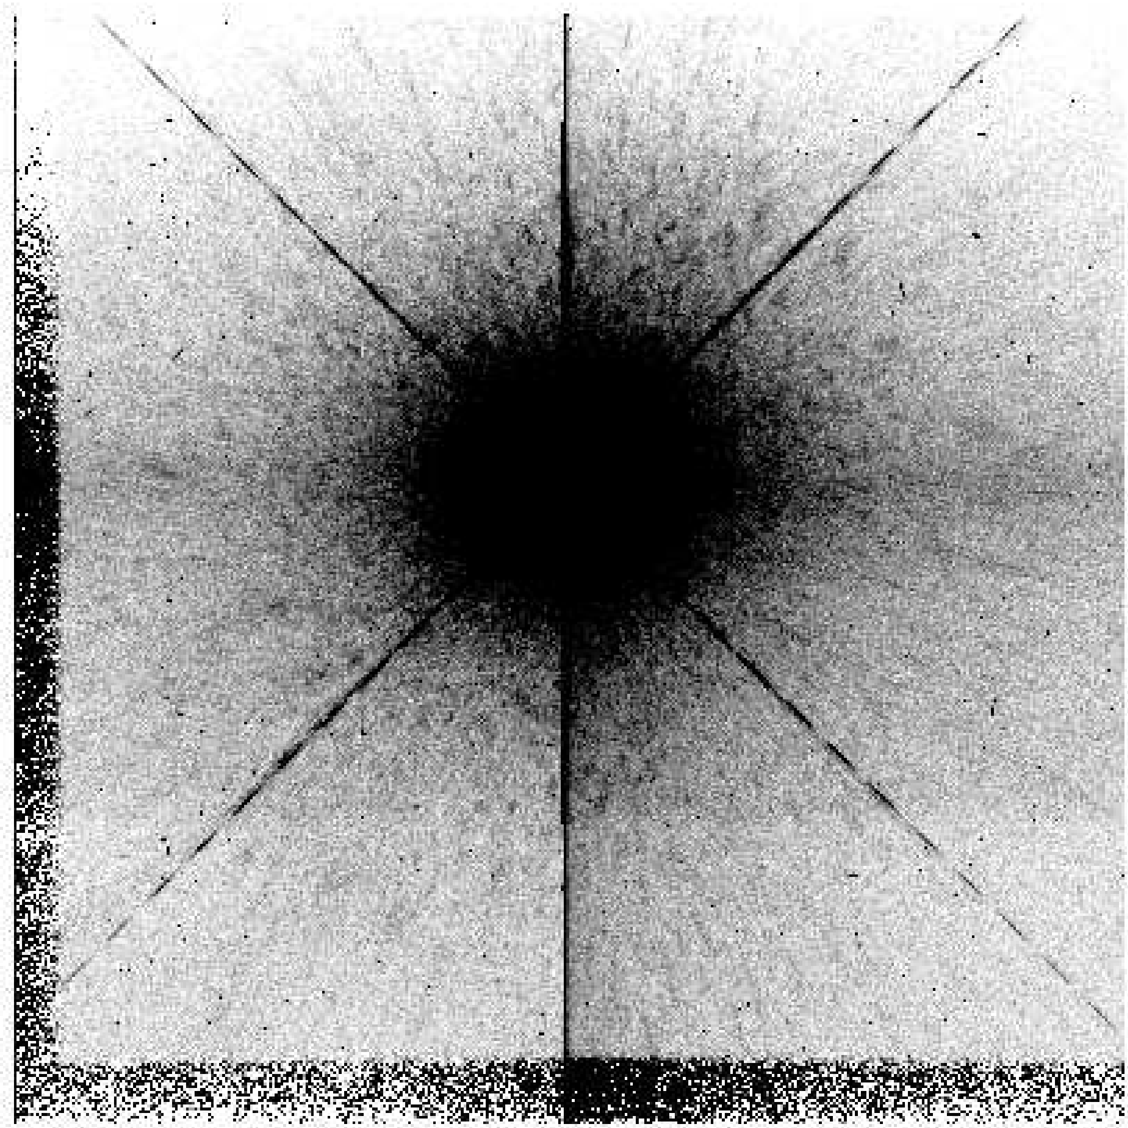

released during an exposure. RBI is temperature sensitive since the bulk trapping time constants decrease with increasing temperature. The WFPC2 CCDs do exhibit RBI, but at -70˚C trapped charge rapidly escapes so that residual images disappear within 1000s. Driven by the WFPC2 electronics, the CCDs recover quickly from large over-exposures (100 times full well or more), showing no measurable residual images a half hour after the overexposure.

For images exposed below the saturation level there is a very weak residual image due to charge trapping and charge transfer efficiency (CTE) problem. Measurements on 1800s dark frames interleaved with 2800s exposures of a star field yield a residual flux of 0.3% ± 0.1% of the original star flux, for stars with fluxes from 65 to 17,000 total counts. For typical star fields observed by WFPC2, these residual images are likely to be a problem only for stars that were saturated in a previous image, or for programs where long exposures in low throughput filters are taken immediately after highly exposes images. Hence, repeated exposures at the same CCD position should not lead to any appreciable systematic offset in photometry. CTE is further discussed in Section 4.11.

4.6 Quantum Efficiency Hysteresis

The problem of quantum efficiency hysteresis (QEH) due to back-side charge accumulation has been reviewed in detail by Griffiths, et al. (1989), and Janesick and Elliot (1991). QEH is not present in the Loral CCDs, because they are front-side illuminated and incorporate MPP operation. This has been verified in component tests at JPL. The absence of QEH means that the devices do not need to be UV-flooded and so decontamination procedures are planned without the constraint of maintaining the UV-flood.

4.7 Flat Field Response

The flat field response is uniform within a few percent, with the exception of a manufacturing pattern defect which generates a 3% reduction in QE once every 34 rows. This pattern defect is identical in all CCDs. It probably creates an astrometric offset of approximately 3% of the pixel height (0.003″ in the WFCs) every 34 rows. More precisely, there was a 0.5μm overlap between adjacent 1024x0.5μm raster scans during the construction of the masks used to fabricate the chips. Photometry of point sources imaged onto these defects will be affected, and it will be better to correct flat fields for these rows for such applications. WFPC2 flat fields also include instrumental effects such as vignetting and shadowing by dust particles.

Figure 4.4: WFPC2 CCD Flat Field

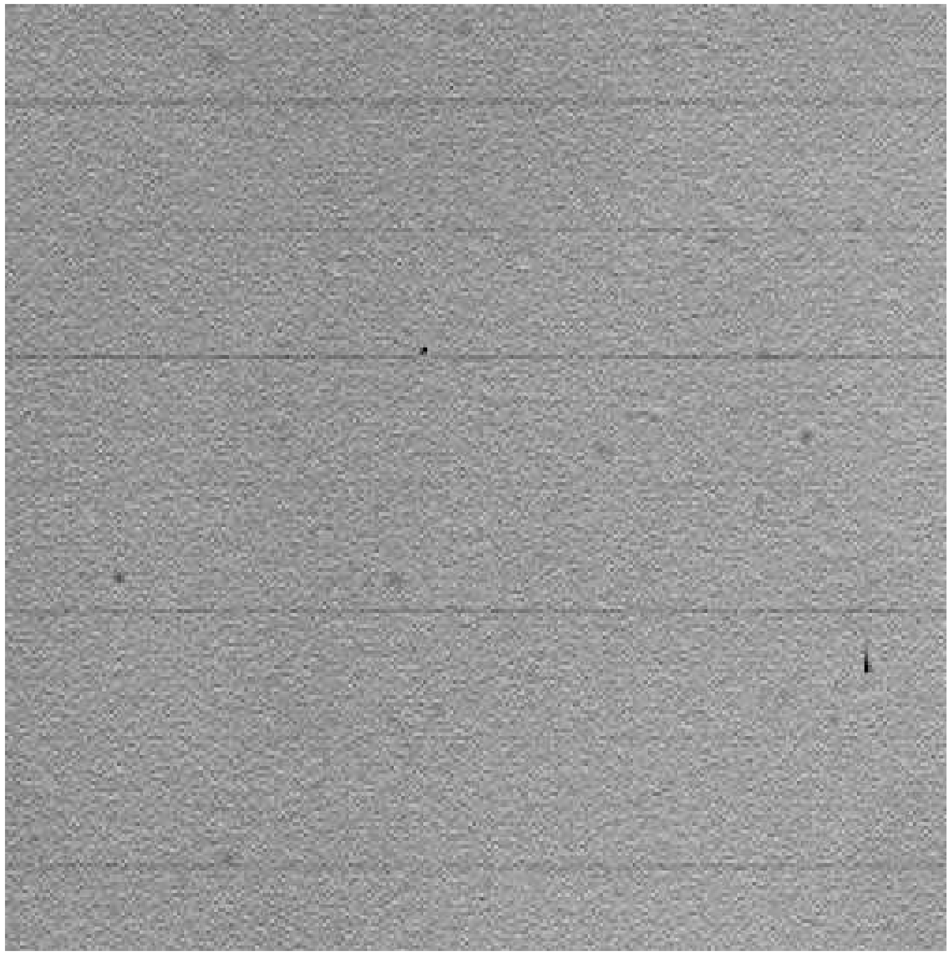

The WFPC2 CCDs have an intrinsically uniform flat field response since they are not thinned, so there are no large-scale chip non-uniformities resulting from the thinning process. MPP operation also improves pixel-to-pixel uniformity because charge transfer is driven deep into the buried n-channel, away from the influence of Si-SiO_2 interface states. The WFPC2 CCD flat fields show an overall pixel-to-pixel response having <2% non-uniformity. Figure 4.4 shows a portion of a WFPC2 CCD flat field obtained during quantum efficiency measurements at JPL. The image illustrates the excellent pixel-to-pixel uniformity of the Loral devices. The 34 row defect is clearly visible, and its amplitude of 3% serves to calibrate the gray scale.

4.8 Dark Backgrounds

Low dark noise is one of the benefits of MPP, since inverted phase operation suppresses the dominant source of CCD dark noise production (Si-SiO_2 surface states). The remaining source of dark noise, thermal generation in the silicon bulk, is determined by the quality of the silicon used in chip fabrication. The intrinsic dark rate of WFPC2 CCDs is <0.01 e^- $pixel^{-1}$ s^{-1} at temperatures below -80°C.

Figure 4.5: Average Dark Rates vs. CCD Row.

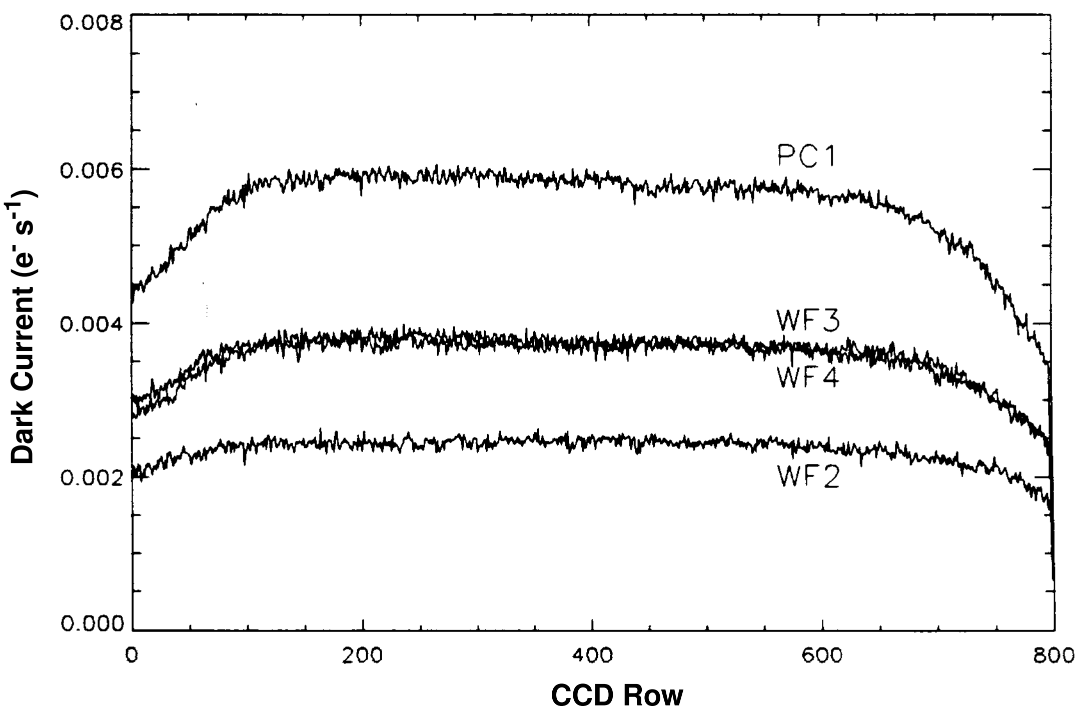

The temperature set-points for the WFPC2 TEC coolers are: -88, -83, -77, -70, -50, -40, -30 and -20 °C. The corresponding approximate median dark rates are given in Table 4.2. For instrument health and safety reasons, GOs cannot command temperature changes.

4.8.1 Sources of Dark Current

The dark current appears to have two components: one from electronic sources in the CCD, and a second component whose strength correlates with the cosmic ray flux. The electronic dark current is ~0.001 e^- s^{-1}, consistent with the Thermal Vacuum Test data.

The second component of dark current appears only on-orbit, its strength drops towards the edges of each CCD, and it is both chip- and time-dependent. At the current operating temperature, this non-electronic component constitutes up to 80% of the total signal measured in the PC. The fraction and overall level are lower in the other chips, and lowest in WF2. This second component ranges from 0.001 e^- s^{-1} (WF2) to 0.005 e^- s^{-1} (PC). The edge drop of is shown in Figure 4.5, where the average of lines 200-600 for each chip (with hot pixels rejected) is plotted in e^- s^{-1} as a function of column number. The drop near the edge is consistent with luminescence from the CCD windows, shadowed by a field stop mask just in front of the CCD.

Table 4.2: Dark Count Rates.

CCD Temperature (˚C)	Dark count rate (e^{-} s^{-1} $pixel^{-1}$)
-20	10.0
-30	3.0
-40	1.0
-50	0.3
-70	0.03
-77	0.016
-83	0.008
-88	0.0045

A further indication of the possible origin of this second component is the correlation between its amplitude and the cosmic ray activity in the same exposure, as shown in Figure 4.6. For example, the cosmic ray flux in the PC varies from 7×10^{5} to 13×10^{5} DN per 1000s, while the total dark signal in the PC varies concurrently between 0.0007 and 0.0010 DN s^{-1}. Similar, though slightly smaller effects are seen in the WFC CCDs. These clues point to cosmic-ray induced scintillation of the MgF_2 field-flattening windows as a likely source of the second dark current component. This might be caused by impurities in the MgF_2 windows; if so, the window of WF2 must contain substantially less impurities. However, other explanations cannot be completely ruled out at this point.

For the great majority of WFPC2 observations, this effect is negligible. In fact, it is noticeable mainly because the true dark rate is very low at the -88˚C operating temperature. However, observations for which the dark current is an important limiting factor, either due to noise or background flatness, will require special handling to remove the signature of the dark current properly, as its amplitude depends on the time-variable cosmic ray flux.

4.8.2 Darktime

As of this writing, the "*DARKTIME*" keyword in the WFPC2 image headers does not reflect correctly the actual time during which the CCD collects dark current. Instead, *DARKTIME* is merely set equal to *EXPTIME* (the exposure time) in the data headers, and this value is used for calibration. The error is small, and usually unimportant, but could be significant for programs aimed at measuring the absolute level of the sky background. The actual darktime in seconds is given by

$$DARKTIME = 60 \times int\left(\frac{t + 13.6}{60}\right) + 46.4 + 13.6 \times (n - 1) + 60 \times (restart) \qquad (4.1)$$

Figure 4.6: Dark Signal vs. Cosmic Ray Flux. Slopes and intercepts ("int") are given on plots. Units are DN/1000s; 1 DN ~ 7 e⁻.

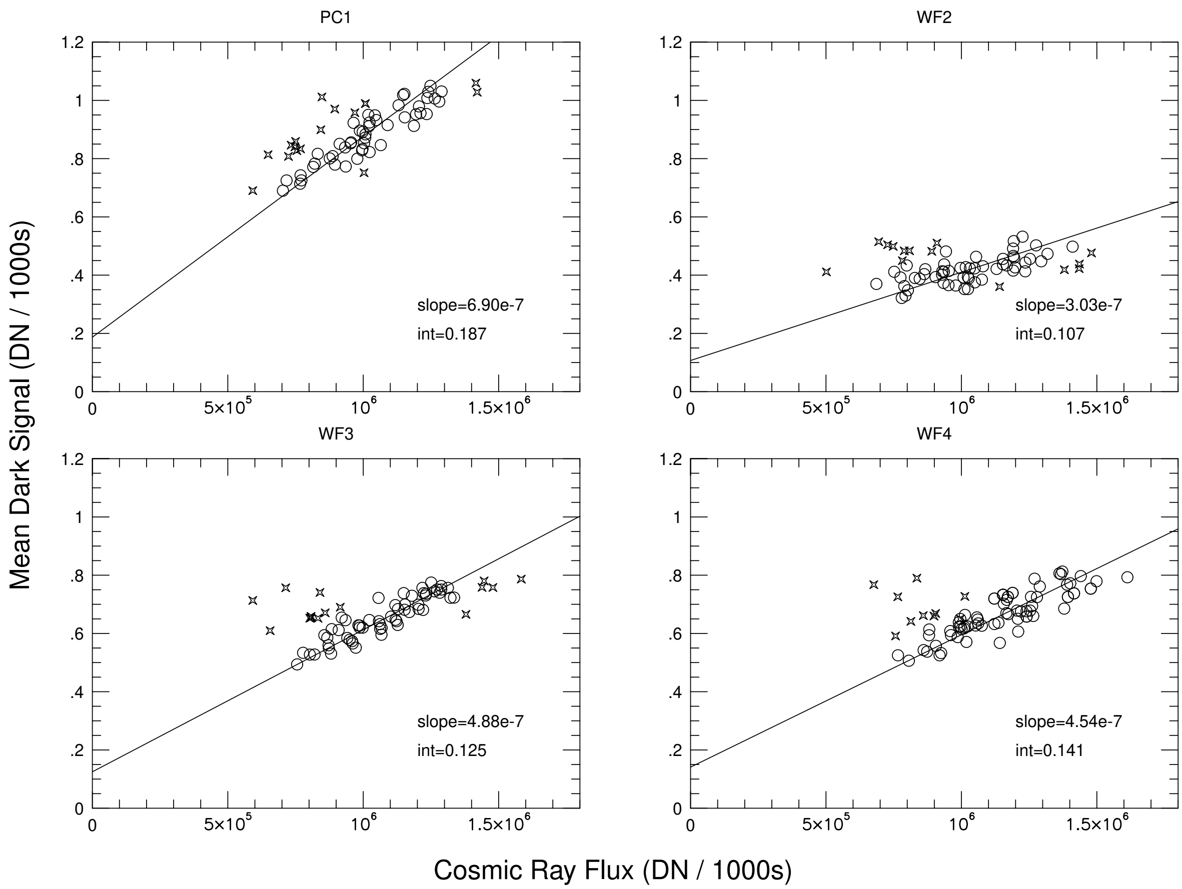

where t is the requested exposure time in seconds, and n is the number of the CCD (PC1=1, WF2=2, etc.), and *int()* indicates the next lower integer. A duration of 13.6s is required both to clear the CCDs before the exposure begins, and to read each CCD after the exposure. External exposures of 180s or longer made with the serial clocks off (CLOCKS=NO; the default setting) suffer an additional 60s of darktime (*restart*=1). This delay is associated with restarting the serial clocks for readout in exposures where the spacecraft AP-17 processor provides shutter control with loss-of-lock checking. Exposures made with the serial clocks on (CLOCKS=YES) avoid this extra 60s (*restart*=0).

We note that bias frames contain approximately $46.4 + 13.6 \times (n - 1)$ seconds of dark current. As of this writing (April 1996), no attempt is made to subtract this from the bias images when creating calibration files for use in the calibration pipeline. This effect is unimportant for most observations, but could be significant if one averaged many undithered deep exposures of the same field, or if one is interested in measuring the absolute level of the sky background. If the dark current were constant in time, this could be corrected by merely changing the value of *DARKTIME* used during calibration. However, the hotpixels vary on monthly timescales, so this simple correction is only partially successful. Work is currently underway to improve the way bias calibration files are computed.

The timing of dark calibration frames is slightly different from that of external science exposures. Dark calibration frames always have *restart*=0 in Equation 4.1.

Dark calibration frames currently used in the calibration pipeline are averages of ten on-orbit dark frames taken over the space of about 2 weeks. The dark calibration file in the pipeline is revised ~every two weeks to track variations in the hotpixels. As of this writing there are plans to alter this method to both reduce the noise and provide better tracking of hotpixels. The new method would use the average of many (~100) dark frames taken over many months, and then hotpixels would be inserted into the calibration file to correct short-term variations. New files would be generated ~weekly, with only the hotpixels changing from week-to-week.

4.9 Cosmic Rays

HST is subjected to cosmic rays and protons from the Earth's radiation belts. The cosmic ray signature in the Loral CCDs is essentially the same as was seen in the WF/PC-1 devices. Electron-hole pairs generated in the thicker substrate by cosmic rays (and infrared photons) are usually removed by recombination in the low resistivity substrate material, because electrons do not diffuse efficiently up to the collecting phase.

Cosmic ray events usually deposit significant quantities of charge in more than one pixel. This is due partly to the finite thickness of the CCD detectors, and partly to the less than perfect collection efficiency of each pixel. Figure 4.7 shows a histogram of the number of affected pixels for each cosmic ray event. For the purposes of the figure, a cosmic ray is defined as having a peak pixel value more than 10 DN above the background; and an affected pixel is an attached pixel with a value more than 2 DN above the background. Cosmic ray events do have a clear lower cutoff at around 500 electrons of total signal. The variations in cosmic ray rates caused by orbital position or operating temperature have not been characterized yet.

Cosmic ray events impact scientific imaging with WFPC2 in two different ways. Firstly, the WFPC2 CCDs have an epitaxial thickness of about 10μm compared to 8μm for the thinned WF/PC-1 device, and a recombination length of 8-10μm in the substrate. These facts lead to a higher total number of electrons being deposited per event. WFPC2 CCDs also have lower read noise, and so the number of cosmic ray events apparently differs from that of the WF/PC-1 CCDs, since low amplitude events are detected. In practice, this means that the number of pixels apparently contaminated by cosmic rays in an image is higher in WFPC2, although the underlying event rate is similar to that experienced in WF/PC-1.

Secondly, stellar images are undersampled and much more difficult to separate from cosmic rays, as is shown in Figure 4.8. Faint stellar images and low energy cosmic rays are indistinguishable. Long science observations are therefore usually

Figure 4.7: Histogram of Cosmic Ray Event Sizes. A cosmic ray event is defined by having a peak pixel of at least 10 DN (at gain 7).

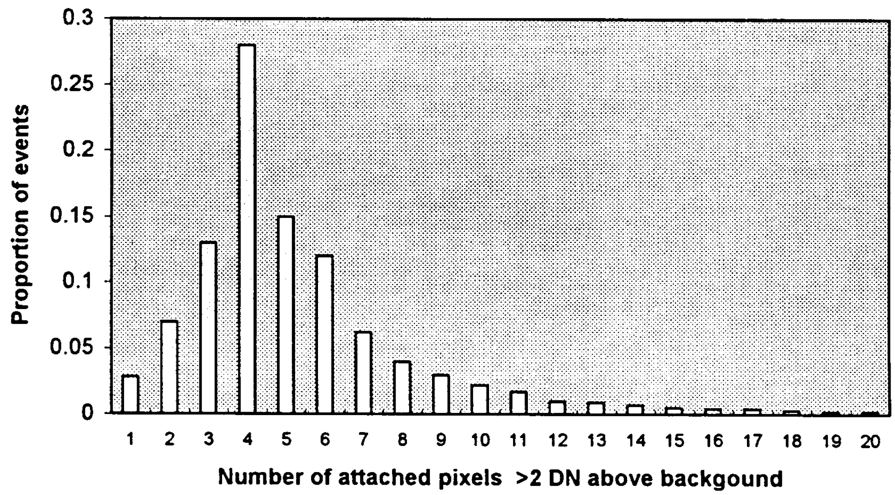

broken into at least two exposures (CR-SPLIT) to ensure that events can be identified.

The average properties of on-orbit cosmic ray events have been determined from examination of several dark exposures, each 2000s long. After bias and dark subtraction, "cosmic rays" were identified in each input frame by first looking for pixels more than 5σ above the background, and then including in each event all adjacent pixels more than 2σ above the background. Very occasionally, two or more physically separate events will be merged into one as a result of this procedure; visual inspection confirms that in the vast majority of cases, this procedure correctly identifies each event and the area affected by it. The typical value of σ was 5 to 6 electrons, including both read and dark noise. The region near the borders of each CCD was excluded in order to avoid edge effects, but all results given here are rescaled to the full area of the CCD.

One difficulty in this measurement is caused by hot pixels, for some of which the dark current has significant fluctuations from frame to frame; these can be mistakenly identified as cosmic rays when the dark current is at a maximum. Single-pixel events constitute 10% of the total number of events identified by our procedure, but at least half of them recur in the same position in several frames, thus identifying them as damaged (hot) pixels, rather than true cosmic rays. Also, unlike the majority of cosmic ray events, single-pixel events tend to have very small total signal; the majority have total signal less than 200 electrons, as expected from hot pixels, while the signal distribution of multiple-pixel events peaks around 1000 electrons. For this reason, single-pixel events have been classified as "bad pixels" rather than "cosmic rays". While we cannot exclude that some true single-pixel events do occur, they are very rare, probably less than 2% of the total.

Figure 4.8: Comparison of Star Images and Cosmic Ray Events. An 80x80 pixel subimage of a 400 second F336W WF2 exposure in ω Cen.

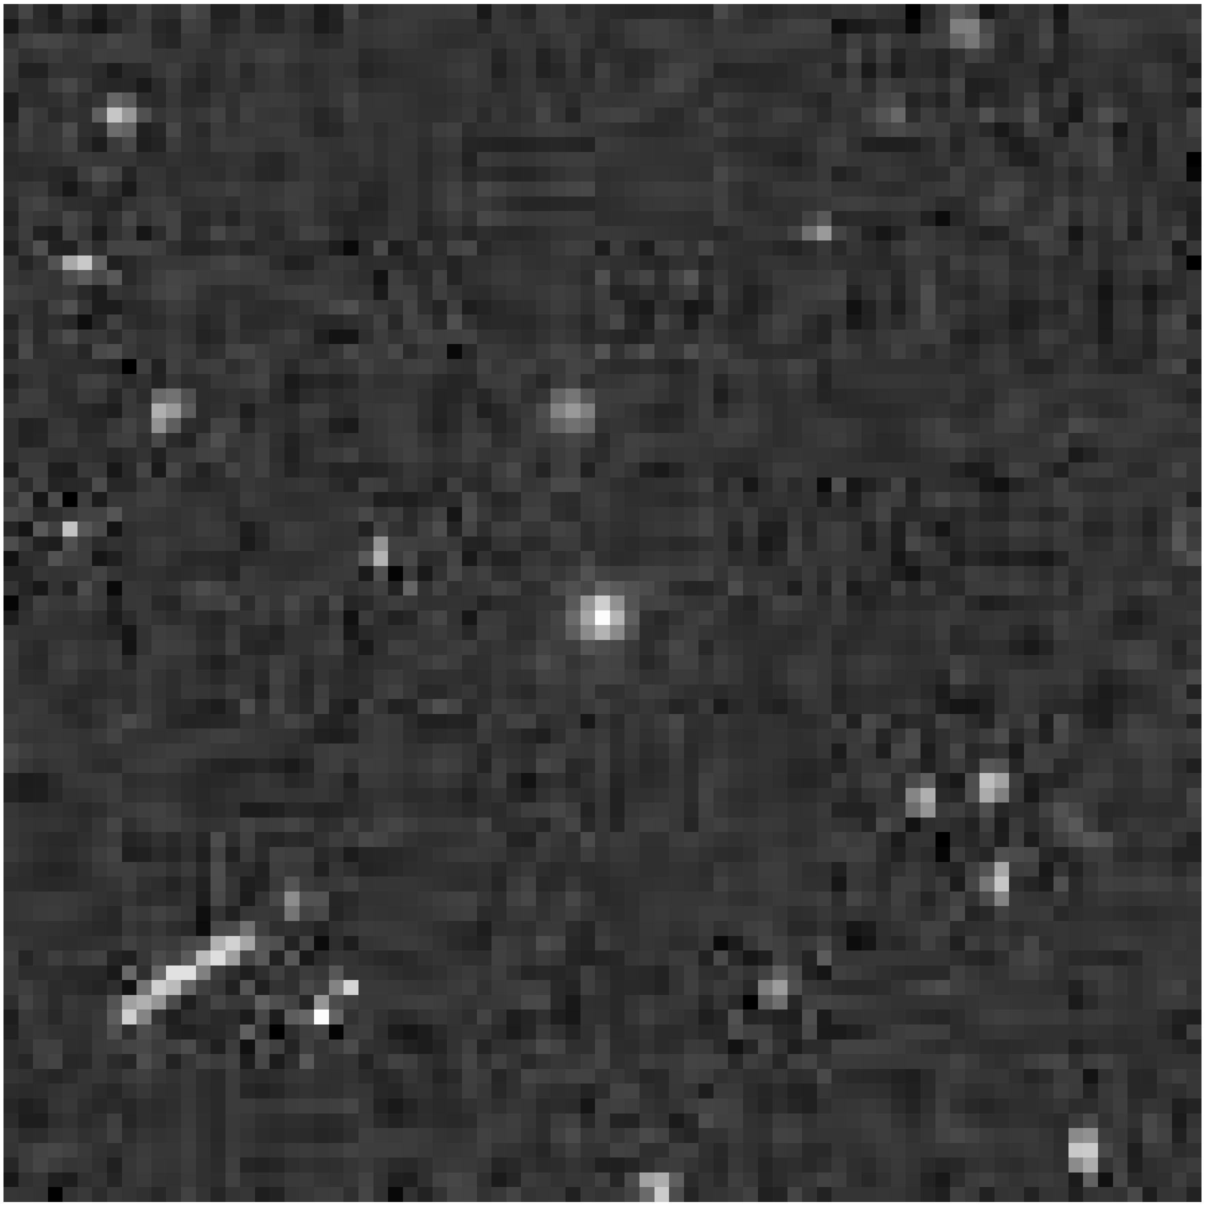

Cosmic ray events are frequent, occurring at an average rate of 1.8 events chip^{-1} s^{-1}. The distribution of total signal is shown in Figure 4.9; it has a well-defined maximum at about 1000 electrons, and a cut-off at about 500 electrons. The latter is well above the detection threshold used for the above measurements (25 electrons in the central pixel of the cosmic ray), and is therefore undoubtedly real.

Figure 4.9: Histogram of Cosmic Ray Event Energies.

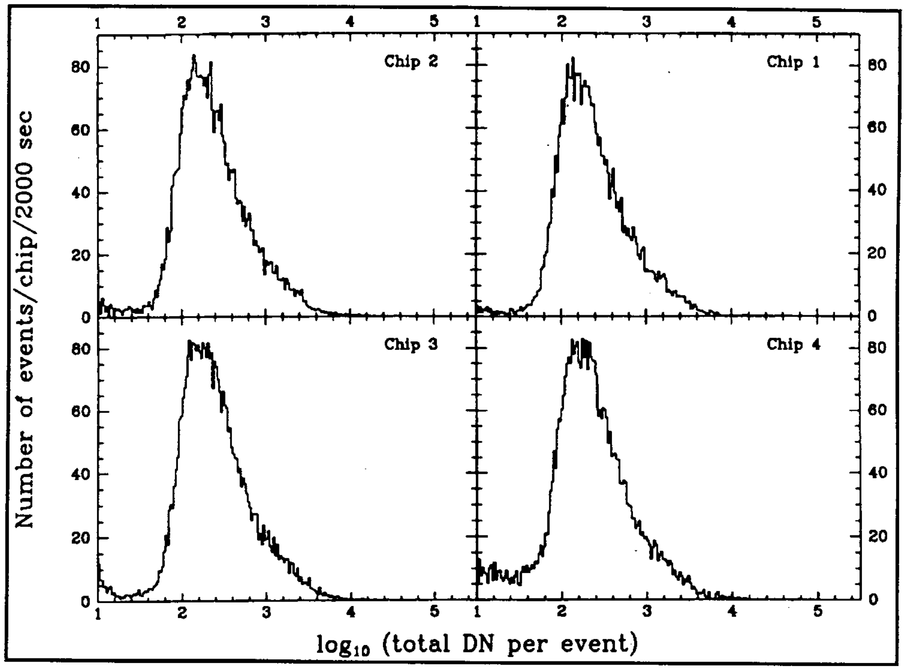

The histogram in Figure 4.9 shows the distribution of total energy of all cosmic ray events. One encouraging feature is the very small number of events below about 30 DN. This low energy drop is well above the energy level of excluded single-pixel events.

A good approximation to the cumulative distribution of events as a function of total signal is given by a Weibull function with exponent 0.25. This function has the form:

$$N(>S) = N_0 \exp[-\lambda(S^{1/4} - S_0^{1/4})]$$

where N is the total number of events which generate a total signal larger than S. The best fit to the observed events gives N_0=1.4 events chip^{-1} s^{-1}, S_0=700 electrons, and $\lambda = 0.57$. The fit fails below S_0, and should not be extrapolated to low-signal events. The rate of events with total signal below 700 electrons is 0.4 events chip^{-1} s^{-1} (i.e. total events per CCD per second is $N_0+0.4 \approx 1.8$).

The number of pixels affected by cosmic ray events in a given exposure is a slightly more sensitive function of the threshold used. While there is a clear drop at low signal for both total and peak signal, neighboring pixels can be affected at low levels. Each event (defined as before) affects an average of 6.7 pixels, for about 12 affected pixels chip^{-1} s^{-1}. For a 2000s exposure, this results in about 24,000 affected pixels, or 3.8% of all pixels. As cosmic rays are expected to be randomly placed, a pair of such exposures would have about 900 pixels affected in both exposures; cosmic ray correction is impossible for such pixels. For a pair of 1000s exposures, about 220 pixels will be affected in both frames.

Cosmic ray activity varies as a function of time, geomagnetic latitude of the spacecraft, and other factors. The average numbers given here are subject to change in individual exposures. After studying about one month's worth of dark exposures, we estimate a total range of about a factor of two in cosmic ray rates.

4.10 Radiation Damage and Hot Pixels

In low Earth orbit (LEO) the CCDs are subject to radiation damage from the Earth's radiation belts. The WFPC CCDs are shielded from energetic electrons and about half the incident energetic protons. Long term radiation damage to the CCDs from high energy protons leads to an increase in dark count rate (mainly from the creation of hot pixels), baseline shifts in the CCD amplifiers, and long term degradation of Charge Transfer Efficiency (CTE). There has not been a significant degradation in the amplifier baselines (although occasional small baseline shifts have been seen) and there have been no observed changes in CTE since launch. On the other hand, one of the major radiation damage mechanisms is the creation of new Si-SiO_2 interface states, which cause increased dark current rates in affected pixels. In the MPP CCD these states immediately recombine with holes, reducing the gradual increase in dark noise by factors of about 25, compared to normal three-phase CCDs (Woodgate, et al., 1989, Janesick, et al., 1989b).

Figure 4.10 is a histogram of the dark current distribution (in $e^- s^{-1}$) for hot pixels. It contains three curves: solid for the histogram of all hot pixels just before a decontamination (April 7, 1995); dashed only for the pixels that were hot just before the decontamination and were not hot at the beginning of the cycle (March 10); and long-dashed for pixels that were hot at the start of the cycle and were fixed by a decontamination. Thus, the dashed line represents the “new” hot pixels, and the long dashed line represents the fixed hot pixels. The fact that these two curves are very similar shows that the number of hot pixels is roughly in equilibrium. The majority of new hot pixels have low dark current. The hot pixels that constitute the accumulated legacy of previous periods, and thus survived one or more decontaminations, include higher-current pixels.

A long term dark frame program is monitoring the increase in the number of hot pixels. The population of hot pixels increases at a rate of approximately 33 pixels CCD^{-1} day^{-1} above a threshold of 0.02 e^- $pixel^{-1}$ s^{-1}, while the camera remains at the normal operating temperature.

About 80% of the new hot pixels return to a normal state at decontamination events when the CCDs are warmed to a temperature of +22°C for 6-12 hours. There is no evidence that the fraction of hot pixels that returns to normal is related to the length of the decontamination. Of those pixels that are not fixed, about half will be fixed after two or three additional decontaminations. After that, the rate of correction decreases. It is conceivable that all hot pixels will be repaired

Figure 4.10: Hot Pixel Histogram.

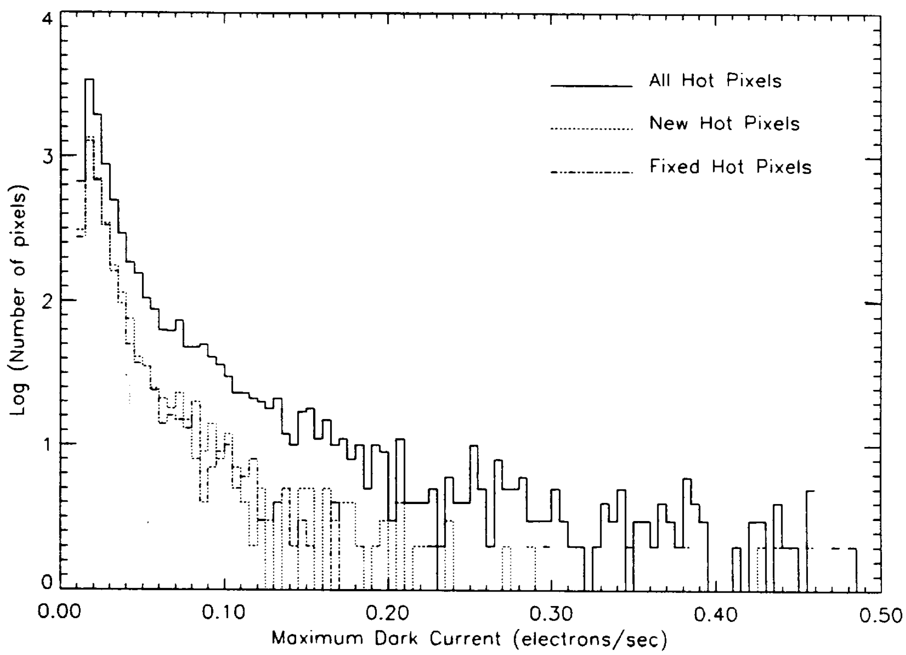

eventually. At the moment there is no evidence of a significant secular increase in the number of hot pixels, and we have a firm upper limit of 8% on the fraction of hot pixels that are not repaired after several decontamination cycles.

In order to deal with the hot pixel problem, we provide monthly lists of possible hot pixels via the World Wide Web and anonymous ftp. Look for hot pixels under WFPC2 Instrument News (/instrument_news/WFPC2/Wfpc2_hotpix via anonymous ftp to stsci.edu). These lists are best used to flag hot pixels as bad. While we do provide an estimate of dark current for each hot pixel as a function of time, there are indications that the noise in hot pixels is much higher than the normal shot noise, and thus dark current subtraction is unlikely to give good results.

4.11 Charge Transfer Efficiency

The WFPC2 CCDs have a small but significant charge transfer efficiency (CTE) problem which causes some signal to be lost when charge is transferred down the chip during readout. This has the effect of making objects at higher row numbers (more charge transfers) appear fainter than they would if they were at low row numbers. The effect depends on the temperature of the CCDs. At the original temperature of -76˚C, as much as 10-15% of the light within a 0.5" radius aperture around a star could be lost for objects at the highest rows. As a result, the CCD operating temperature was changed to -88˚C on 23 April 1994. Now the effect seems to have a maximum amplitude of 4% for stars with more than 1,500 total detected electrons and no measurable difference in slope is seen for stars up to 20,000 total electrons. However, the brightest stars seem to have a smaller fractional loss. It is also possible that proportional losses for fainter stars, at least with no background, may be larger. Hence, the effect is not well described by either a constant fractional or a constant additive loss per charge transfer. The effect also depends on the amount of background on the chip. There is significantly less CTE effect in the presence of even a moderate (several hundred electrons) background. A major calibration effort is underway to quantify better the dependencies on star color and brightness, and background (see Chapter 8).

The photometric calibration of the instrument presented here (and this summary of CTE) is based on Holtzman, et al. (1995b). It has been crudely corrected for CTE. All of the frames considered are short exposures with essentially no background. For data taken at -88˚C, a 4% correction ramp was applied to the measured 0.5″ radius aperture photometry, in the sense that objects at row 800 were made brighter by 4%, but the brightness of objects at the first row was not changed. The correction was applied to bring measurements to the values they would have had in the absence of CTE, or equivalently, the values they would have had if measurements had been made at row 0. Thus, the calibration presented here applies directly to images in which CTE is insignificant (for example, faint stars with high background levels). For images with low background, a CTE correction must be applied before using the calibration presented in this handbook.

The current lack of a comprehensive understanding of CTE effects introduces one of the largest uncertainties for WFPC2 photometry. The CTE problems are caused by electron traps in the CCDs which are filled as charge passes through pixels. However, not all traps are accessible to all electrons passing through. Some traps are only accessible if there is significant charge involved. This model suggests that there will not be significant CTE losses in the presence of background, particularly for faint stars, because background electrons fill the traps before such stars pass through. For brighter stars with background there will still be some loss because their charge may access traps that are unaffected by the background that previously clocked through. Faint stars in scenes with little background may suffer from larger losses, although there is no direct evidence for this yet.

An apparently related effect surfaced during Cycle 5 as observers compared long and short exposures of the several stellar fields -- there was evidence for a discrepancy between the photometric zeropoint of long and short exposures. Further investigation shows that “long vs. short” is probably a misnomer. The level of background appears to be the main parameter rather than the exposure time. The magnitude difference measured between short and long exposures is more pronounced for faint stars in large apertures, where it can reach 0.05 mag, and is essentially absent for stars with more than 1000 total counts. The dependence on aperture and magnitude appears consistent with a charge transfer efficiency problem. The offset of faint star magnitudes can be explained by a loss of 0.3 DN (2 e^-) in each pixel used in the aperture. Further testing and analysis of this effect are underway.

A side effect of the CTE problem, is that residual trails can appear in exposures after a highly exposed image is read out (Figure 4.11). The trail is caused by charge which is trapped during read out of the highly exposed image, which is then slowly released during subsequent exposures. The effect is most pronounced when long exposures in through low throughput filters (narrow band or UV filters) immediately follow a highly exposed image (usually a broad band filter).

4.12 Read Noise and Gain Settings

The CCDs and their associated signal chains have readout noise levels (in the absence of signal shot noise or interference) of approximately 5e^-. The analog-to-digital converter is highly accurate, and makes virtually no contribution to the read noise, other than the normal information loss caused by digitization of the signal.

The conversion factors from detected electrons (QE x number of incident photons) to data numbers (DN) are tabulated in Table 4.3, as are read noise and linearity ("gamma" is the power law index relating detected DN to input flux). Note that all calculations of sensitivity in this manual assume gains of 7 and 14 for two gain channels, choices very close to the measured gains. The photometric calibration is based on an assumed exact gain of 14 in all CCDs. The measurements given here were derived from thermal vacuum testing. On-orbit measurements have confirmed that the gain ratios are correct to within a possible systematic error of 1%—which will feed directly into a photometric calibration error for gain 7 data, as most of the photometric calibration was done with gain 14 data. Note that the gain ratios are known much more accurately than the individual gains; they are derived from flat field ratios instead. Also, note that the Phase II proposal instructions refer the ~14 e^- DN^{-1} setting as ATD-GAIN=15.

Table 4.3: Signal Chain Gains.

Parameter	Gain	PC1	WF2	WF3	WF4
Noise	"7"	5.2 4 ± 0.30	5.51 ± 0.37	5.22 ± 0.28	5.19 ± 0.36
	"15"	7.02 ± 0.41	7.84 ± 0.46	6.99 ± 0.38	8.32 ± 0.46
Gain	"7"	7.12 ± 0.41	7.12 ± 0.41	6.90 ± 0.32	7.10 ± 0.39
	"15"	13.99 ± 0.63	14.50 ± 0.77	13.95 ± 0.63	13.95 ± 0.70
Gamma	"7"	1.0015 ± 0.0006	1.0015 ± 0.0006	1.0020 ± 0.0006	1.0038 ± 0.0007
	"15"	1.0004 ± 0.0001	1.0023 ± 0.0004	1.0032 ± 0.0006	1.0018 ± 0.0012
14/7	ratio	1.987 ± 0.02	2.003 ± 0.02	2.006 ± 0.02	1.955 ± 0.02

Figure 4.11: Images Illustrating CTE Residual Trail. (a) Image of star field taken in broad band filter on PC1. (b) 1500s dark exposure taken immediately following (a). Read out direction is towards bottom of image. Cosmic rays have been removed.

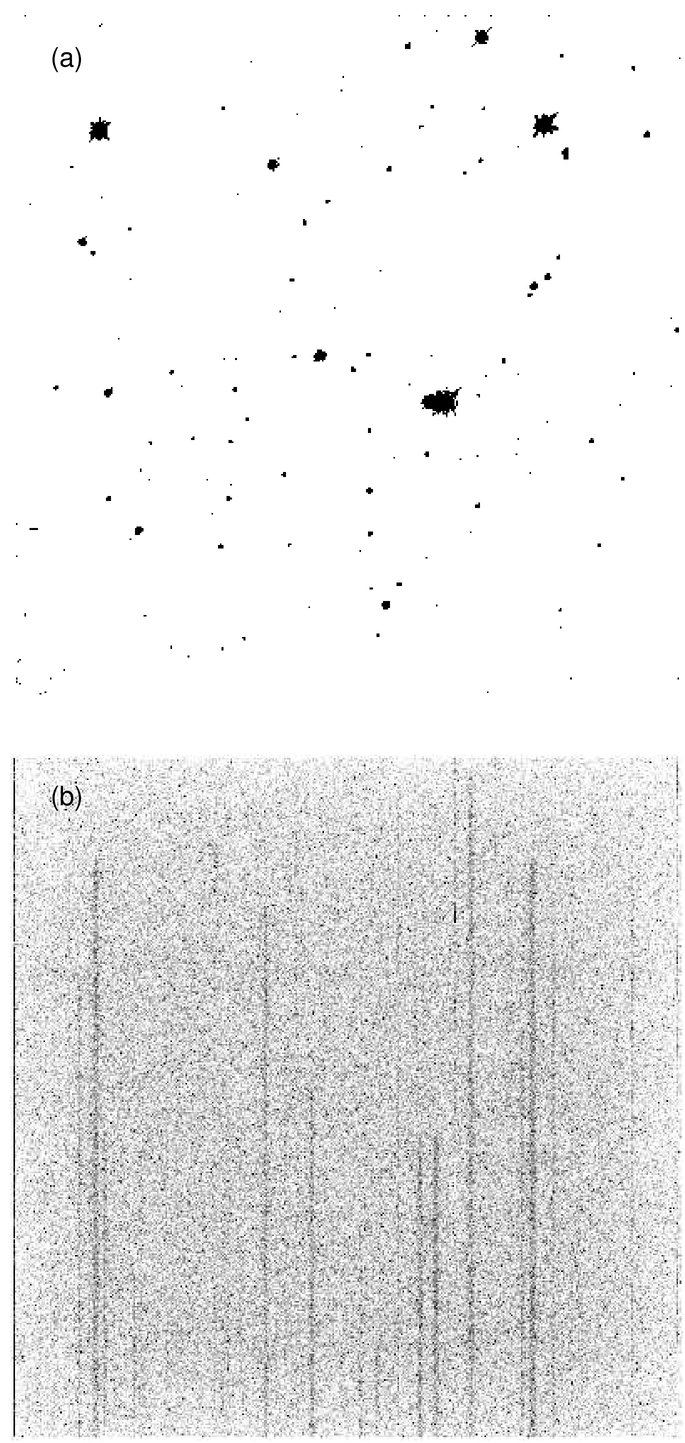

CHAPTER 5:

Point Spread Function

In This Chapter...

5.1 Effects of OTA Spherical Aberration

The OTA spherical aberration produces a Point Spread Function (PSF—the apparent surface brightness profile of a point source), as presented to the instruments, with broad wings. Briefly, the fraction of the light within the central 0.1″ was reduced by a factor of about 5. The resulting PSF had "wings" which extended to large radii (several arcseconds), greatly reducing the contrast of the images and degrading the measurements of sources near bright objects or in crowded fields. Burrows et al. (1991, Ap. J. Lett. 369, L21) provide a more complete description of the aberrated HST PSF. Figure 5.1 shows the PSF in three cases. It shows the aberrated HST PSF, the WFPC2 PSF, and for comparison the PSF that would be obtained from a long integration if HST were installed at a ground based observatory with one arcsecond seeing. All of the PSFs were computed at 4000Å. The FWHM of the image both before and after the installation of WFPC2 is approximately proportional to wavelength, at least before detector resolution and MTF effects are considered. (The WF/PC-1 core was approximately 50% broader than the core that is obtained with WFPC2). Figure 5.2 shows the encircled energy (EE), the proportion of the total energy from a point source within a given radius of the image center, for the same three

Figure 5.1: PSF Surface Brightness. The percentage of the total flux at 4000Å falling on a PC pixel as a function of the distance from the peak of a star image.

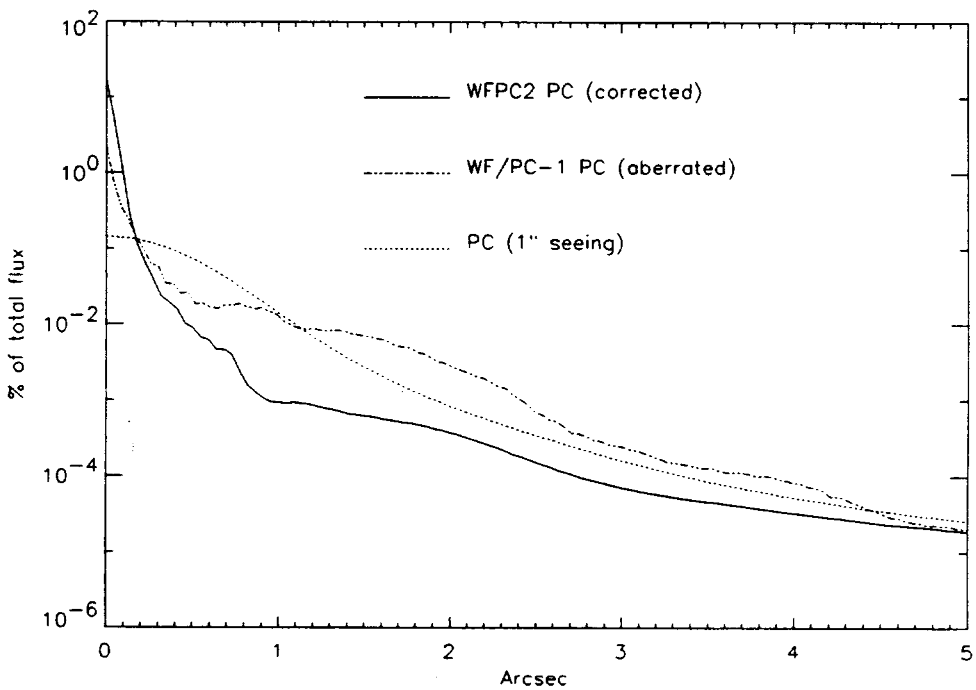

cases. The WFPC2 curve shown is the average of measurements taken with F336W and F439W. It can be seen that the core of the image in WFPC2 contains most of the light. At this wavelength, 65% of the light is contained within a circle of radius 0.1″. However, this proportion is considerably less than the optics deliver. The reason for this is discussed in the Section 5.4. Encircled energy curves for other filters are shown in Figures 5.3 and 5.4; note that these curves are normalized to unity at 1.0″ radius.

5.2 Aberration Correction

WFPC2 has corrective figures on the relay secondary mirrors where the primary mirror is imaged; this optical correction recovers near-diffraction limited images over the entire CCD fields-of-view. Proper correction requires tight optical alignment tolerances, which are facilitated on-orbit by actuated optics. The corrective optics enable essentially all of the scientific objectives of the original WF/PC-1 to be met.

Figure 5.2: Encircled Energy. The percentage of the total flux at 4000Å within a given radius of the image peak.

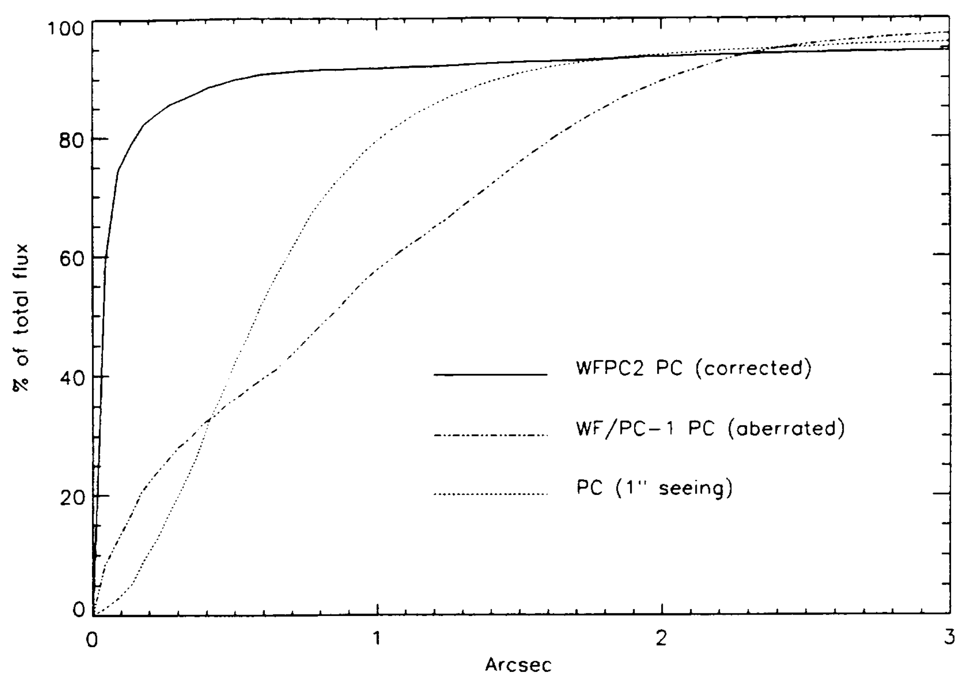

Table 5.1: Wavefront Error Budget.

Camera	WFC (F/12.9)	PC(F/28.3)
Design error	λ/143	λ/50
Fabrication and alignment error	λ/14.7	λ/14.7
Alignment stability error	λ/25	λ/25
Total wavefront error	λ/12.6	λ/12.3

Through a number of independent analyses, based on investigations of star images obtained on-orbit, and the examination of fixtures used during the figuring of the primary mirror, the aberrations of the HST optics were accurately characterized. The primary mirror was figured to an incorrect conic constant: –1.0139±0.005 rather than the –1.0023 design requirement, resulting in a large amount of spherical aberration. The optical design of WFPC2 creates an image of the OTA primary mirror near the surface of the relay Cassegrain secondary mirror in each of its channels. This design minimizes vignetting in the relay optics, but more importantly, facilitates correction of spherical aberration in the OTA primary by application of the same error (but with opposite sign) to the relay secondary. The optical figure of the WFPC2 secondary mirrors contains a compensating

Figure 5.3: Encircled Energy for CCD PC1. The fraction of energy encircled is plotted vs. aperture radius for several filters. Curves are normalized to unity at a radius of 1.0″. From Holtzman, et al., 1995a.

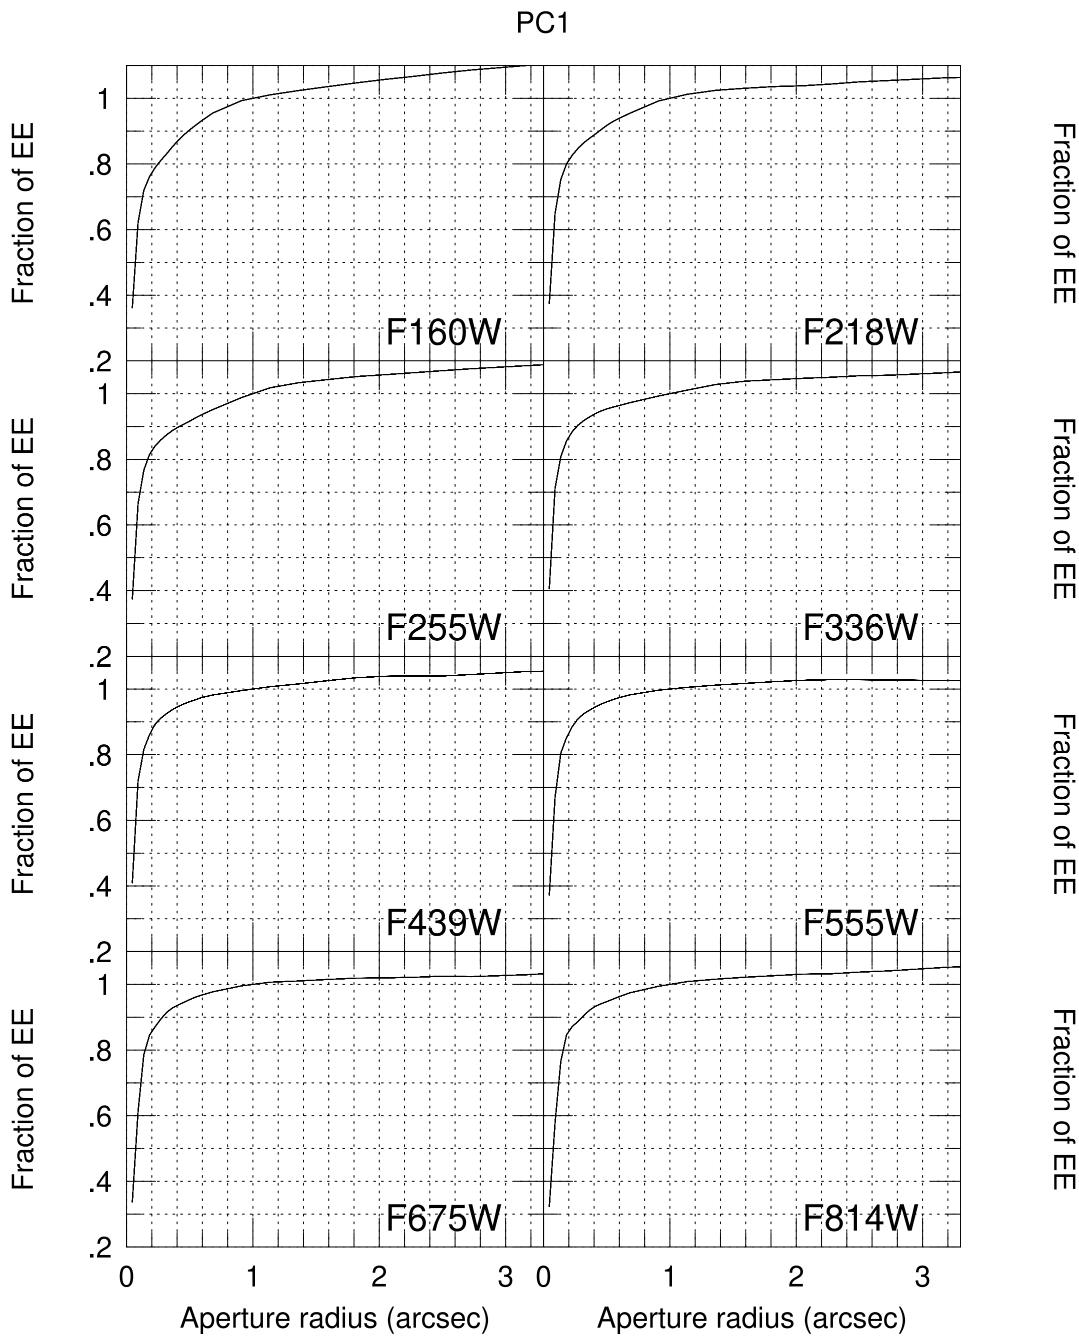

Figure 5.4: Encircled Energy for CCD WF3. The fraction of energy encircled is plotted vs. aperture radius for several filters. Curves are normalized to unity at a radius of 1.0″. From Holtzman, et al., 1995a.

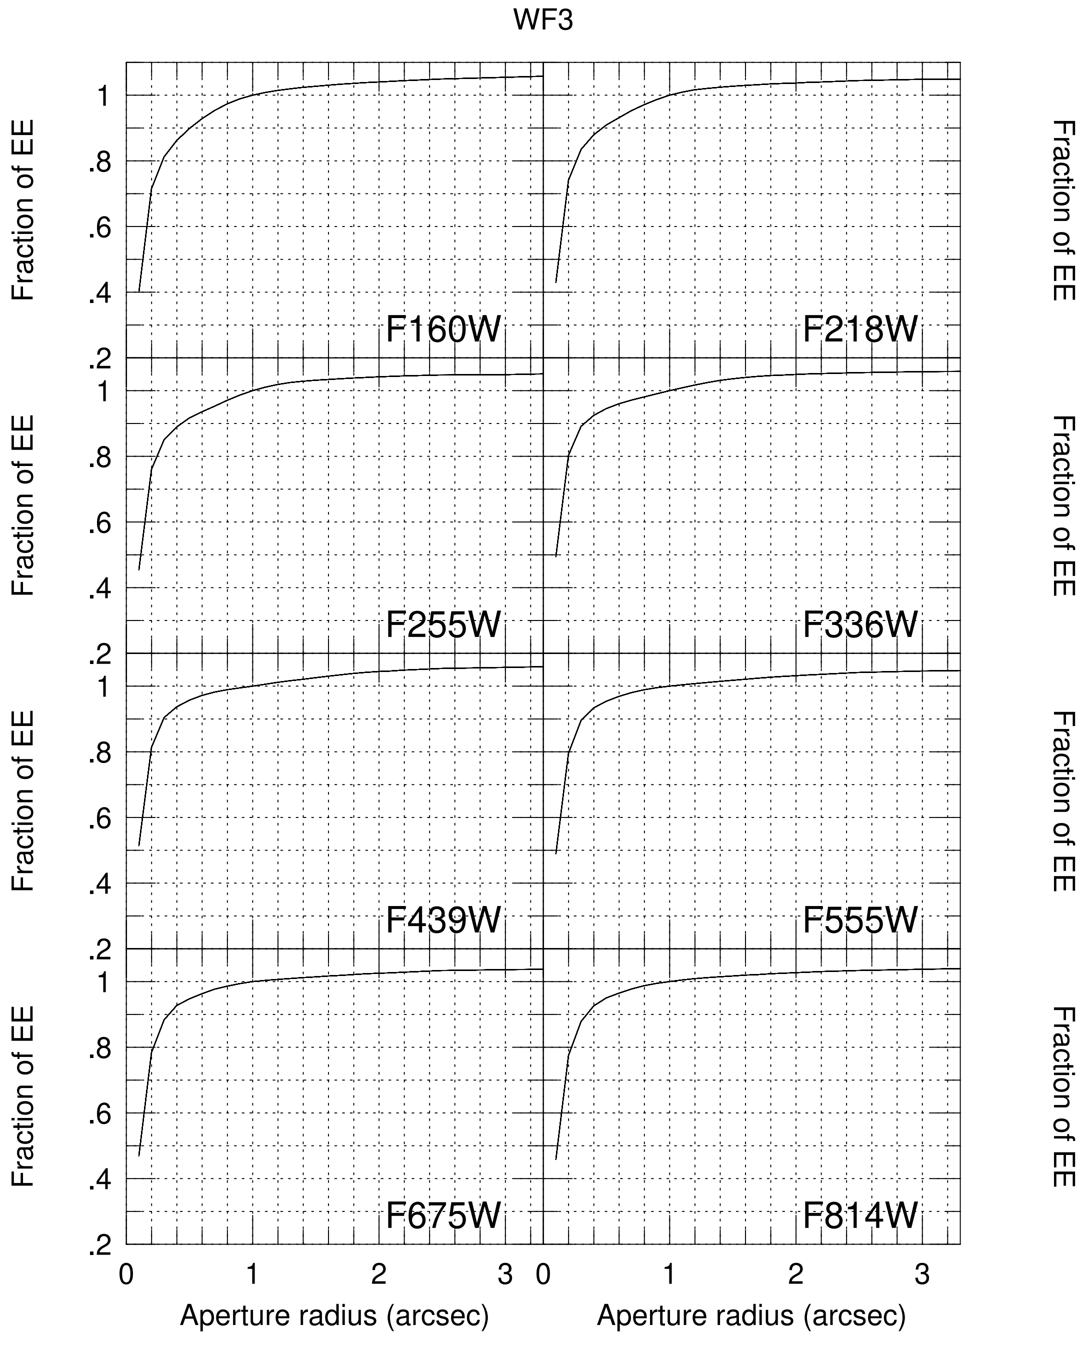

"error" in the conic constant. By adopting a prescription within the error bars for the HST primary mirror, corrective secondary mirrors were made with sufficient accuracy that the residual spherical aberration in the WFPC2 wavefront is small compared to other terms in the WFPC2 optical wavefront budget.

On the other hand, new and stringent alignment requirements were created by the steep optical figure on the corrective relay secondary mirrors. The primary mirror image must be accurately centered on the corrective mirror, and must have the correct magnification. Centering is the most demanding requirement. A failure to center accurately would create a new aberration in the form of coma. A misalignment of 7% of the pupil diameter introduces as much RMS wavefront error as was present in the form of spherical aberration prior to the introduction of corrective optics. The new requirements for alignment accuracy and stability led to the introduction of a tip-tilt mechanism on the pick-off mirror, to compensate for camera alignment uncertainties with respect to the OTA, and actuated fold mirrors which can compensate for internal misalignments. There was an additional term in the CEIS specification of the overall instrument wavefront error budget for alignment stability. It is $\lambda/25$ RMS at 6328Å, as shown in Table 5.1.

"Design error" refers to the aberrations inherent in the design itself, which would be seen if the optics conformed perfectly to their specifications. All of the optics were fabricated and integrated into the WFPC2 optical bench. It was established on the basis of component tests, end-to-end optical interferometry, and through focus phase retrieval, that the WFPC2 optical system performed within the stated tolerances for "fabrication and alignment" in the laboratory environment. What remained was to demonstrate the stability of the optical alignment after launch vibration and in response to the thermal environment on-orbit. The "stability" line anticipated these uncertainties, and has been verified during early science operations.

5.3 Wavefront Quality

The conclusion of the extensive optical testing in thermal vacuum was that the cameras are well corrected to within the specifications. The measured wavefront errors in the four cameras were 1/30, 1/17, 1/40, and 1/21 waves at 6328Å. The dominant problem was a small difference in focus between the four cameras. The actuated fold mirrors and pick-off mirror mechanism performed flawlessly in correcting residual coma aberrations in the image, and enabled the on-orbit alignment procedures. By defocusing the OTA secondary mirror, a very accurate alignment of the cameras was accomplished. A side effect was that the aberrations in each camera were measured (Krist and Burrows, Applied Optics, 1995). The results are given in Table 5.2. These values were used in generating the simulated PSFs given in section 5.5, "Model PSFs", on page 92.

Table 5.2: Aberrations in Each Camera. The numbers quoted are RMS wavefront errors in microns over the HST aperture (Zernike coefficients).

	Aberration	PC1	WF2	WF3	WF4
Z4	Defocus	0.0000	0.0410	0.0640	0.0480
Z5	0° Astig	0.0229	0.0109	0.0126	0.0163
Z6	45° Astig	0.0105	0.0041	0.0113	0.0190
Z7	V2 Coma	0.0000	0.0012	-0.0037	-0.0090
Z8	V3 Coma	-0.0082	0.0061	-0.0100	0.0019
Z9	X Clover	0.0063	0.0121	0.0010	0.0096
Z10	Y Clover	0.0023	0.0091	0.0130	0.0042
Z11	Spherical	-0.0131	-0.0215	-0.0265	-0.0247
Z22	5th Spherical	0.0034	0.0034	0.0036	0.0029
	Zonal Errors	0.0180	0.0180	0.0180	0.0180
	Total (RS5)	0.0353	0.0537	0.0755	0.0637
	Budget	0.0813	0.0794	0.0794	0.0794

5.4 CCD Pixel Response Function

From Thermal Vacuum testing, there was evidence that the images are not as sharp as expected, despite the good wavefront quality. The decrease in sharpness corresponds to a loss in limiting magnitude of about 0.5 magnitudes in the WF cameras, and less in the PC.

Further testing, by covering a flight spare CCD with a 2µm pinhole grid in an opaque metallic mask and illuminating it with a flat field source, showed that even when a pinhole was centered over a pixel only about 70% of the light was detected in that pixel.

For practical purposes, the effect can be modeled as equivalent to about 40 mas RMS gaussian jitter in the WFC, and 18 mas in the PC (as compared with the typical real pointing jitter of ~3 mas delivered by the excellent HST pointing control system). Alternatively, at least in the V band, it can be modeled by convolving a simulated image by the following kernel, which gives the pixel response function averaged within pixels:

$$K = \begin{bmatrix} 0.016 & 0.067 & 0.016 \\ 0.080 & 0.635 & 0.080 \\ 0.015 & 0.078 & 0.015 \end{bmatrix}$$

One clue is the wavelength dependence of the observed sharpness: the results from the 2µm pinhole grid test get worse at longer wavelengths. This may reflect the greater penetration into the silicon of low energy photons, which facilitates the diffusion of photoelectrons across the pixel boundaries defined by the frontside gate structure.

There is also evidence for sub-pixel QE variations at the 10% level. There is an implied dependence on pixel phase for stellar photometry. This has been seen at about the 1% level in on-orbit data. The work of Jorden, Deltorn, and Oates (Greenwich Observatory Newsletter 9/93) has yielded quite similar results, and suggests that sub-pixel response must be taken into account when seeking to understand the behavior of all CCD detectors forming undersampled images.

5.5 Model PSFs

Considerable effort has gone into the modeling of the HST point spread function (PSF), both in order to measure the optical aberrations in support of the WFPC2, COSTAR, and advanced scientific instruments, and to provide PSFs for image deconvolution in the aberrated telescope. Such PSFs are noise free and do not require valuable HST observing time. Software to generate model PSFs for any filter and at any location within the field-of-view is available from the STScI (TIM package, Hasan and Burrows 1993; TinyTIM package, Krist 1995). The results are illustrated in Tables 5.3 and 5.4 for the PC1 and WF2 cameras, respectively. A representative PSF is on the left in each panel. It meets the wavefront error budget, with the measured mix of focus, coma, astigmatism, and spherical aberration. It has been degraded by the pixel response function as discussed in the Section 5.4. On the right is the diffraction limited case for comparison. In each case the percentage of the total flux in a central 5x5 pixel region of a point source is displayed. The peak of the star image can be at an arbitrary point relative to the boundaries of the CCD pixels. Two cases are shown: one where the star is approximately centered on a pixel, and one where it is approximately centered at a pixel corner. As a consequence of the under-sampling in the WFPC2, the limiting magnitude attainable in the background limit varies by about 0.5 magnitude, depending on the position of the source within the CCD pixel. This point is discussed in more detail in Chapter 6.

Neither observed nor modeled PSFs will provide a perfect match to the PSF in actual science observations, due to modeling uncertainties, the “jitter” in the HST pointing, and orbit to orbit variations in telescope focus (“breathing”—which seems to be generally limited to about 1/20 wave peak-to-peak). Jitter is not predictable but can be recovered to a reasonable extent for observations obtained in Fine Lock. In long exposures, up to about 10 mas of apparent pointing drift may occur as a result of the breathing effects in the FGS, although smaller variations of ~3 mas are typical.

Table 5.3: PC Point Spread Functions. Shown as percentages (out of 100 percent) of the total flux in a 5 by 5 pixel region. On the left in each case is a model PSF with the observed wavefront errors and pixel response function. On the right is the diffraction limited case for comparison.

WFPC2 Model PSF					2000 Å: Peak near corner of PC pixel			Diffraction Limited PSF	
0.9	2.3	2.3	0.9	0.3	0.2	0.6	0.5	0.3	0.1
2.5	10.3	12.7	2.5	0.5	0.6	17.6	20.9	0.5	0.1
1.9	11.2	13.3	2.6	0.5	0.5	20.9	26.0	0.6	0.2
0.9	2.1	2.7	1.4	0.3	0.3	0.5	0.6	0.3	0.1
0.3	0.5	0.5	0.4	0.2	0.1	0.1	0.2	0.1	0.2
					Peak near center of pixel				
0.3	0.7	1.3	0.8	0.5	0.1	0.3	0.4	0.2	0.2
0.7	2.6	6.4	3.2	0.8	0.3	0.6	4.9	0.6	0.3
1.2	6.3	25.0	6.9	1.4	0.4	4.9	62.9	6.3	0.4
0.6	2.2	5.9	3.6	0.9	0.3	0.5	6.3	0.7	0.4
0.4	0.8	1.4	1.1	0.4	0.2	0.3	0.4	0.4	0.2
					4000 Å: Peak near corner of PC pixel				
0.9	2.5	3.5	1.2	0.2	0.3	2.5	2.5	0.3	0.1
3.6	10.8	12.0	3.1	0.4	2.5	14.3	15.8	2.7	0.2
2.7	11.5	12.9	3.5	0.4	2.5	15.8	17.6	3.0	0.2
0.8	3.0	3.4	1.3	0.3	0.3	2.7	3.0	0.4	0.1
0.2	0.4	0.4	0.3	0.2	0.1	0.2	0.2	0.1	0.1
					Peak near center of pixel				
0.3	0.6	0.8	0.9	0.3	0.1	0.2	0.3	0.2	0.1
0.9	4.0	6.6	4.9	0.8	0.2	3.8	4.4	3.8	0.2
0.9	6.9	26.1	7.0	0.8	0.3	4.5	49.4	5.1	0.4
0.5	3.4	6.8	4.9	0.9	0.2	3.8	5.0	4.3	0.2
0.2	0.6	0.8	0.9	0.4	0.1	0.2	0.4	0.2	0.1
					6000 Å: Peak near corner of PC pixel				
2.0	2.6	3.4	2.4	0.5	2.1	2.3	2.0	2.1	0.2
3.4	9.8	10.4	2.9	0.6	2.2	11.4	12.8	2.1	0.6
2.8	10.6	11.2	3.2	0.6	2.0	12.9	14.1	2.1	0.6
1.6	2.8	3.1	2.4	0.5	2.0	2.0	2.1	2.3	0.2
0.4	0.6	0.7	0.5	0.2	0.2	0.5	0.6	0.2	0.1
					Peak near center of pixel				
0.5	1.2	1.7	1.7	0.6	0.2	1.5	1.8	1.4	0.3
1.7	3.1	5.9	3.6	1.6	1.5	2.4	4.3	2.2	1.6
2.0	6.0	20.7	6.6	1.8	1.8	4.4	31.6	5.3	2.0
1.2	2.9	6.2	3.7	1.7	1.4	2.1	5.4	2.3	1.7
0.4	1.3	2.0	1.7	0.6	0.2	1.5	1.9	1.7	0.3
					8000 Å: Peak near corner of PC pixel				
1.6	1.9	2.2	1.9	1.1	1.8	0.9	0.9	1.5	1.1
2.1	9.3	9.7	2.1	1.1	0.9	11.7	12.6	1.0	1.4
2.0	9.8	10.1	2.4	1.1	0.9	12.6	13.3	1.0	1.5
1.4	2.1	2.1	1.8	1.1	1.5	1.0	1.0	1.6	1.2
0.8	1.2	1.3	1.1	0.4	1.1	1.4	1.5	1.2	0.2
					Peak near center of pixel				
1.2	1.4	1.5	1.8	1.4	1.3	1.8	1.1	1.6	1.4
1.8	2.5	6.0	2.8	1.6	1.8	1.5	6.2	1.7	1.6
1.6	6.0	15.4	6.6	1.5	1.1	6.2	22.5	7.1	1.0
1.3	2.7	6.3	3.0	1.7	1.6	1.7	7.1	1.9	1.7
1.0	1.4	1.5	1.7	1.4	1.4	1.6	1.0	1.7	1.5

Table 5.4: WFC Point Spread Functions.Shown as percentages (out of 100 percent) of the total flux in a 5 by 5 pixel region. On the left in each case is a model PSF with the observed wavefront errors and pixel response function. On the right is the diffraction limited case for comparison.

WFPC2 Model PSF					2000 Å: Peak near corner of WF pixel		Diffraction Limited PSF		
0.5	1.8	2.1	0.8	0.3	0.4	0.4	0.4	0.3	0.1
1.5	9.9	13.6	3.1	0.5	0.4	15.4	21.5	0.4	0.1
1.5	9.8	24.0	4.7	0.5	0.4	21.5	33.4	0.4	0.1
0.5	2.1	3.6	1.3	0.3	0.3	0.4	0.4	0.4	0.1
0.2	0.4	0.4	0.2	0.2	0.1	0.1	0.1	0.1	0.1
					Peak near center of pixel				
0.2	0.4	0.5	0.5	0.3	0.2	0.2	0.2	0.1	0.2
0.3	1.7	5.3	3.0	0.8	0.1	0.8	1.3	0.6	0.1
0.5	5.4	28.5	14.0	2.2	0.2	1.3	86.2	1.4	0.2
0.3	2.2	9.1	4.5	1.0	0.1	0.6	1.4	0.9	0.2
0.3	0.6	1.2	0.8	0.3	0.2	0.1	0.2	0.2	0.2
					4000 Å: Peak near corner of WF pixel				
0.8	2.6	2.5	0.8	0.2	0.4	0.5	0.5	0.2	0.1
2.6	16.1	16.4	3.0	0.4	0.5	17.7	21.2	0.6	0.1
2.0	12.9	17.3	3.0	0.4	0.5	21.2	26.7	0.6	0.1
0.6	2.0	2.6	0.9	0.2	0.2	0.6	0.6	0.4	0.2
0.2	0.4	0.3	0.2	0.1	0.1	0.1	0.1	0.2	0.2
					Peak near center of pixel				
0.3	0.6	0.7	0.5	0.3	0.3	0.2	0.3	0.2	0.2
0.6	3.1	7.2	3.4	0.7	0.2	0.7	3.7	0.7	0.2
0.9	8.5	33.3	10.2	1.2	0.3	3.7	68.8	5.3	0.4
0.4	2.5	8.8	3.0	0.6	0.2	0.7	5.3	0.8	0.2
0.2	0.5	0.9	0.5	0.3	0.2	0.2	0.4	0.2	0.3
					6000 Å: Peak near corner of WF pixel				
0.7	2.6	2.6	0.9	0.3	0.2	0.5	0.4	0.2	0.2
3.0	14.9	15.6	3.3	0.5	0.5	18.3	20.7	0.5	0.3
2.2	13.7	16.3	3.2	0.4	0.4	20.7	24.2	0.6	0.2
0.6	2.3	2.9	0.7	0.2	0.2	0.5	0.6	0.2	0.2
0.2	0.3	0.3	0.2	0.2	0.2	0.3	0.2	0.2	0.2
					Peak near center of pixel				
0.3	0.7	0.9	0.6	0.2	0.2	0.3	0.2	0.3	0.2
0.7	4.1	7.8	4.1	0.8	0.3	1.8	6.1	1.9	0.3
1.0	8.6	30.4	9.4	1.3	0.2	6.1	54.9	6.2	0.3
0.5	2.8	8.2	3.5	0.6	0.3	1.9	6.2	2.5	0.3
0.2	0.5	1.0	0.6	0.3	0.2	0.3	0.3	0.3	0.2
					8000 Å: Peak near corner of WF pixel				
1.0	3.0	2.9	1.1	0.3	0.1	2.0	2.0	0.2	0.2
3.6	13.1	13.6	4.0	0.5	2.0	15.8	17.3	2.2	0.1
2.6	12.6	14.2	3.5	0.5	2.0	17.3	19.3	2.5	0.1
0.8	3.2	3.6	1.0	0.3	0.2	2.2	2.5	0.2	0.2
0.2	0.4	0.5	0.2	0.1	0.2	0.1	0.1	0.2	0.1
					Peak near center of pixel				
0.2	0.8	0.8	0.7	0.3	0.1	0.2	0.3	0.2	0.1
0.9	4.6	6.9	4.4	1.0	0.2	3.5	4.6	3.4	0.2
0.9	7.5	30.8	8.1	1.2	0.3	4.6	52.0	5.0	0.3
0.5	3.1	7.3	3.7	0.7	0.1	3.4	5.1	3.9	0.2
0.2	0.5	1.0	0.7	0.2	0.1	0.2	0.3	0.2	0.1

5.6 PSF Variations with Field Position

The WFPC2 PSFs vary with field position due to field-dependent aberrations, obscuration shifting, and scattering. This will complicate photometry, PSF subtraction, and deconvolution (Krist, 1995).

The coma and astigmatism aberrations vary significantly within a camera across the field-of-view. These variations are simply part of the optical design. At the extreme corners of the WFC CCDs, away from the OTA axis, there is about 1/5 wave of astigmatism (referenced at 633 nm), which decreases to nearly zero at the CCD centers. Astigmatism at this level causes the PSF core to become elliptical and slightly less sharp; note flattening of PSF at pixels positions (54,777) and (605,148) in Figure 5.5. Coma also varies, but to a much lesser extent. Coma and astigmatism variations are considerably smaller in PC1.

Figure 5.5: PSF Variations with Field Position - Aberrations. Nine observed PSFs (filter F814W) are shown from a widely spaced grid on WF3. CCD pixel position are labeled. Note flattening of PSF in the (54,777) and (605,148) positions.

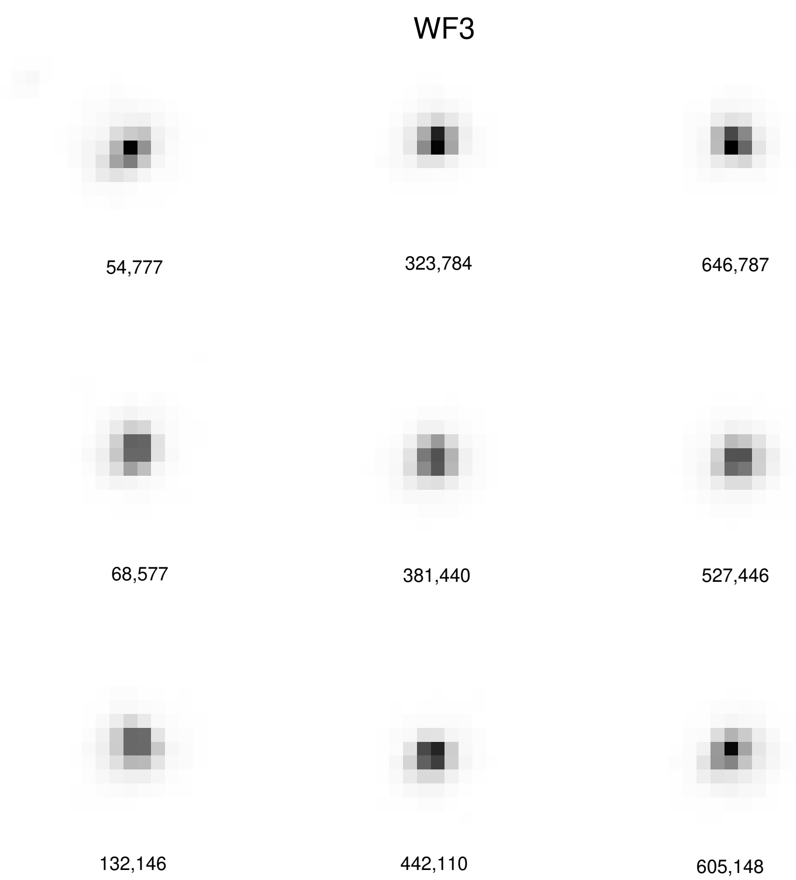

The obscuration patterns due to the camera optics (relay secondary mirror and spiders) appear to shift with respect to the OTA obscurations, depending on field position. The interacting diffraction patterns of the WFPC2 and OTA spiders cause ripples in the spider diffraction spikes, which vary with field position as the two spiders shift relative to each other. In Figure 5.6 the OTA spider is hidden behind the WFPC2 spider at the field center and hence the diffraction spikes there have a simple, smooth appearance (c.f. position 446,425). At the CCD corners, however, one or more vanes of the OTA spider moves out from behind the WFPC2 spider, and the double set of obscurations causes a "beating" pattern in the diffraction spikes.

The spiders also interact with light diffracted from zonal errors in the OTA mirrors, causing streaks in the scattering halo which vary in position and intensity.

Figure 5.6: PSF Variations with Field Position - Obscuration Shifts. Five saturated PSFs observed in F814W are shown from a widely spaced grid on WF4. Note changes in spider diffraction spikes. CCD pixel position are labeled.

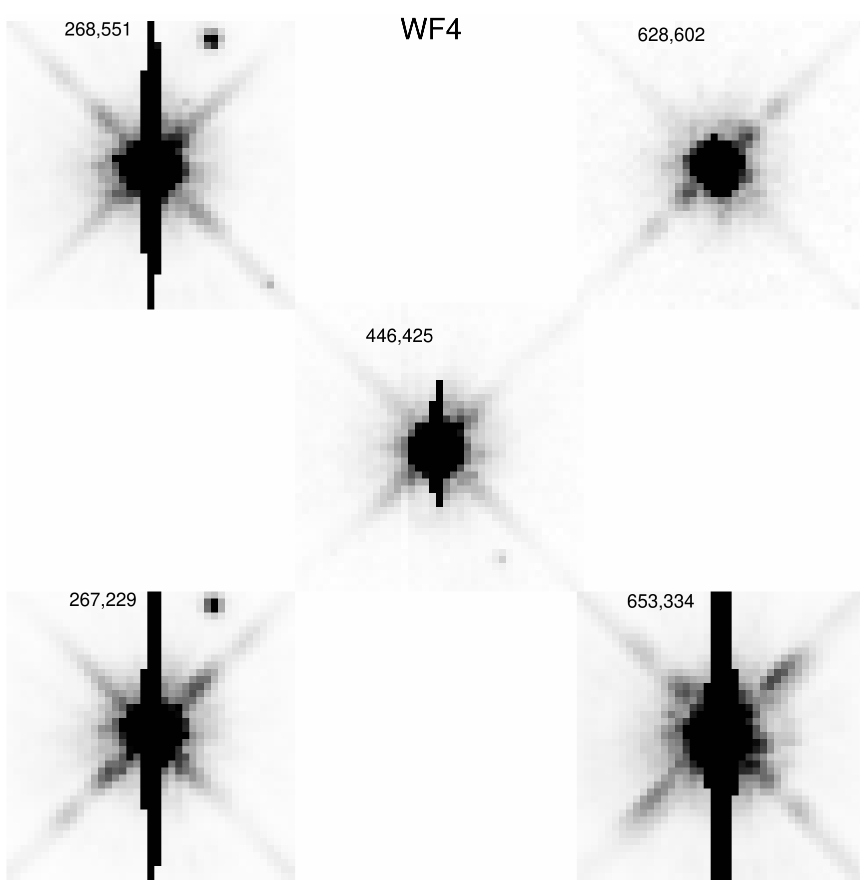

5.7 PSF Variations with Time

The shape and width of observed PSFs varies slightly over time, due to the change in focus of the telescope. The focus variation consists of two terms: a secular change, due to the ongoing shrinkage of the Metering Truss Assembly, at a currently estimated rate of 0.85 µm month^{-1}, and short-term variations, typically on an orbital time-scale (the so-called "breathing" of the telescope). The breathing is probably due to changes in the thermal environment as the telescope moves through its orbit, and has a typical peak-to-peak amplitude of 4 µm; larger variations are occasionally seen.

These small focus shifts will impact photometry performed with small (few pixel radius) apertures. Typical ±2µm focus shifts will result in photometric variations in the PC1 of 6.8%, 4.5%, 2.0%, and 0.2% for aperture radii of 1, 2, 3, and 5 pixels, respectively in F555W. This is based on the focus monitoring data taken over the period from July 1994 to January 1996 (Figure 5.7). Hence, "breathing" is often one of the major sources of errors for small-aperture photometry. However, relative photometry (i.e. the difference in magnitudes of stars in the same image) is not affected by this variation very much since all the stars in an image are impacted by the defocusing in a similar way. The variations are, of course, smaller for the WF chips because the pixels are roughly twice as large.

Figure 5.7: Measured OTA Focus Position (microns) as Function of Number of Days since January 1, 1994. Focus position is defined as difference between the optimal position of OTA secondary mirror and actual position, in microns.

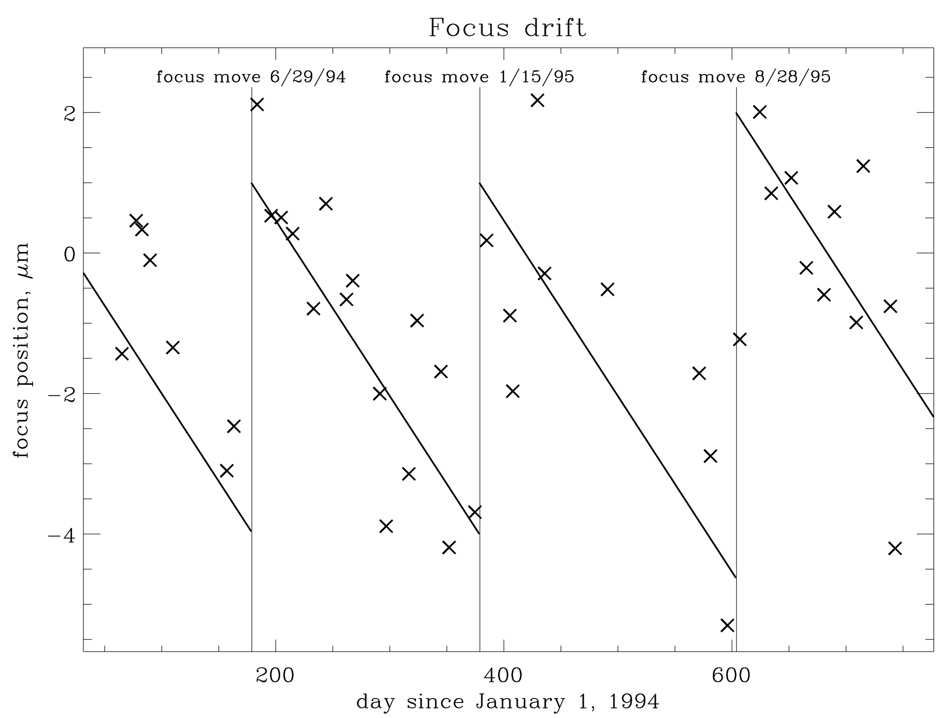

Figure 5.8: Measured Aperture Correction, V(r) - V(r=10 pix), in Magnitudes as Function of Shift from Optimal Focus. Data are given for aperture radii r=1, 2, 3, and 5 pixels for F555W filter on CCD PC1.

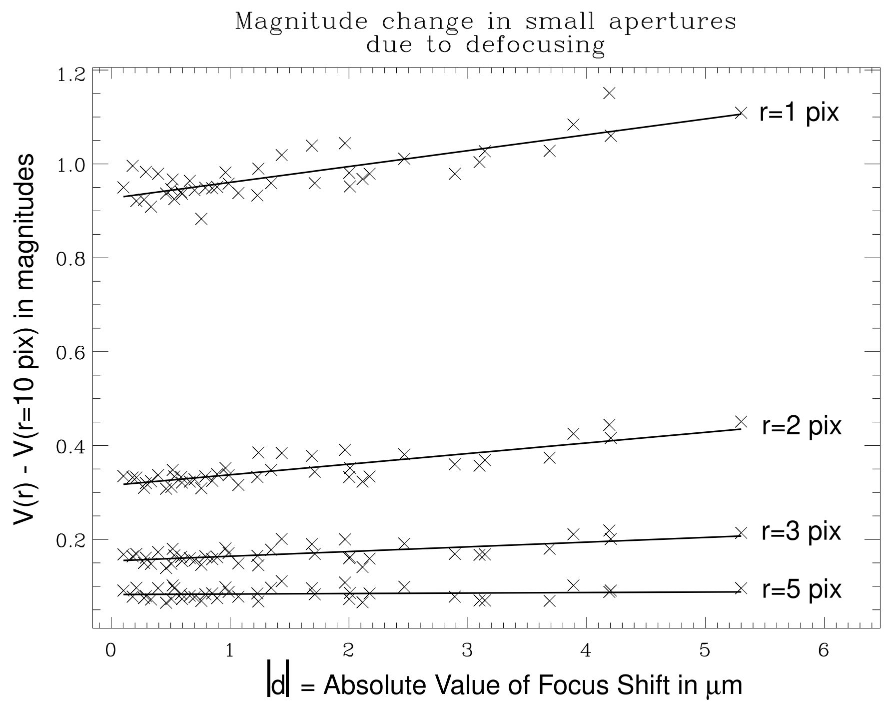

Systematic errors due to the secular focus drift can be corrected using the aperture corrections phased with the focus change (see Figure 5.8):

ap_corr = ap_corr_nominal + a(r) x d

where ap_corr_nominal is the nominal aperture correction (mag) as derived from Table 2a in Holtzman et al. (1995a), a(r) is the flux variation per 1μm of focus drift (mag per micron) using an aperture with radius r (pixels), and d (μm) is the focus shift from the nominal position. The monitoring data mentioned above yield for PC1, F555W, the following values for a(r):

a(1 pix) = 0.0338 ± 0.0038

a(2 pix) = 0.0226 ± 0.0024

a(3 pix) = 0.0100 ± 0.0018

a(5 pix) = 0.00105 ± 0.0015

Values for other filters and other chips will be compiled in an instrument science report later in 1996.

Large focus changes, with amplitudes up to 10μm, are seen occasionally (See Hasan and Bely, Restoration of HST Images and Spectra II, p. 157). On May 1, 1994, and February 27, 1995, a short-lived defocusing of the telescope of up to 10μm was seen, probably due to extreme thermal conditions. Such a defocusing causes an increase of the PSF width by about 5-10% and a significant change in its shape, especially evident in the PC both because of its higher resolution and its astigmatism (which makes the out-of-focus image appear elongated). The change in the PSF appears to be modeled adequately by the TIM software. (See Hasan and Bely, Restoration of HST Images and Spectra II, p. 157). (See sample PSF subtraction in Figure 7.2 on page 148).

5.8 Large Angle Scattering

Analysis of the WFPC2 saturated star images indicate that the large angle scattering (>3″ from a star) is significantly higher than expected.

Three data sets were used to determine the WFPC2 scattering. The first set was from the SMOV Ghost Check proposal 5615, in which 100 second images of δ Cas (V=2.7) were obtained at the center of each chip in F502N. The second set was a series of 6 second exposures of Vega (V=0.0) centered on WF2 through F410M (WFPC2 GTO proposal 5205). The third set was ε Eridani (V=3.73) centered on the PC and taken through F631N (500 sec each) and F953N (2200 sec each). These were from GTO proposal 5611.

WFPC2 scattering was determined by computing the azimuthal average and azimuthal median profiles. The regions near the diffraction spikes and saturated columns were not used. The profiles were determined using images corrected for horizontal smearing. After renormalizing to the XCAL fluxes the profiles agreed.

The measurements indicate that the average scatter in WFPC2 is an order of magnitude greater than in WF/PC- 1. The increase is due to scattering in WFPC2, not the OTA. In the WFPC2 images, the pyramid edge shadow is not visible in the scattered light; the light is spread out to the chip edges, indicating that most of the scattering occurs after the pyramid. However, the light level in adjacent channels is back down at the WF/PC-1 levels as shown in Figure 5.9.

The scattering does not show any strong dependence on wavelength between 410 nm and 953 nm, within the uncertainties of the measurements.

The scattered light is not uniform. There are high frequency spatial structures in the form of streaks radiating outwards from the star. These features are probably both wavelength and position dependent, and so cannot be readily subtracted.

The source of the WFPC2 scattering may be the CCDs. The WF/PC-l CCDs were back illuminated and had shiny surfaces. The electrode structure was not visible over most of the wavelength range. The WFPC2 CCDs, however, are front illuminated, so the electrode structure is visible and may be scattering the light. There was a large ghost in WF/PC-l due to a reflection between the CCD and filter, but no such feature has been seen in WFPC2. The flux from this missing ghost may instead constitute part of the scatter. (See also related material in “Observing Faint Targets Near Bright Objects” on page 146.)

Figure 5.9: Large Angle Scattering. The proportion of the total flux in F555W falling per square arcsecond as a function of the distance from the peak of a saturated stellar image. These curves are for a target in the PC. Note the large drop in the scattered light level when looking in an adjacent camera.

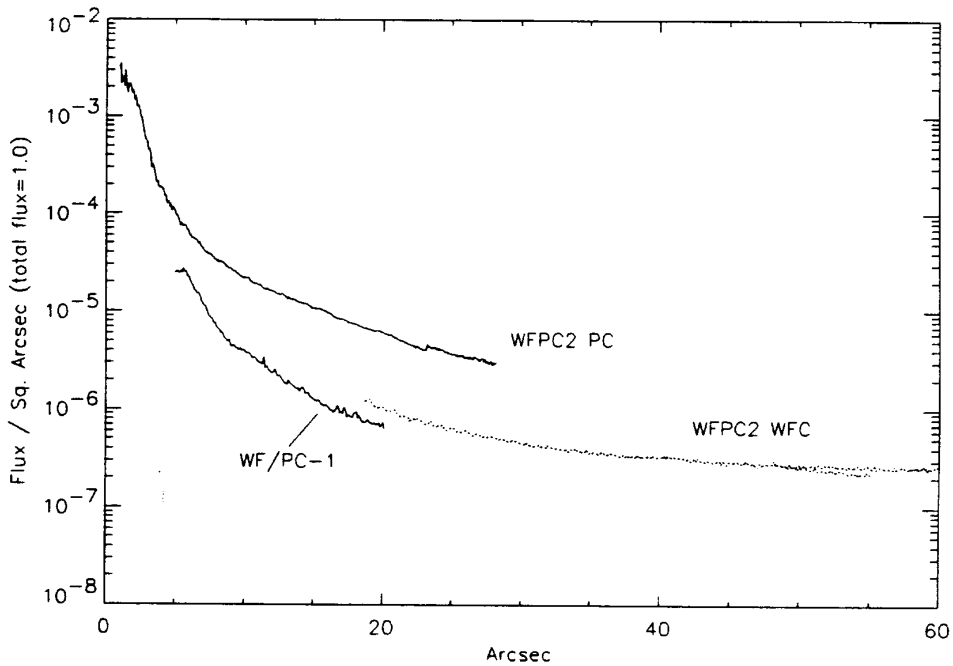

5.9 Ghost Images

Common ghost images result from internal reflections in the filters and in the field-flatteners. Two filter ghosts, caused by double (and quadruple) reflection inside the filter, are visible below and to the right of the star in Figure 5.10. The position and brightness of these ghosts varies from filter to filter, typically being most obvious in interference filters. The comatic shape of the ghost is caused by the camera optics being effectively misaligned for the light path followed by the ghost. The relative position of these ghosts does not vary much over the field.

An additional ghost is caused by an internal reflection inside the MgF_2 field flattener lens immediately in front of each CCD (Figure 5.11). The field flattener ghost is doughnut shaped (image of OTA pupil) in the WFC, but is smaller and more disk-like on the PC. This ghost contains ~0.15% of the total energy of the star. It is positioned on a line through the CCD center and the bright star; the distance from the ghost to the CCD center is 1.25 to 1.4 times the distance from the bright star to the CCD center. This geometry results from curvature of the field flattener lens.

The large ghost image expected to be caused by reflection off the CCD back to the filter and then back to the CCD is not seen. It was deliberately eliminated in the PC by tilting the CCD slightly.

Figure 5.10: Saturated Stellar Image Showing Filter Ghosts. Intensity scale is logarithmic.

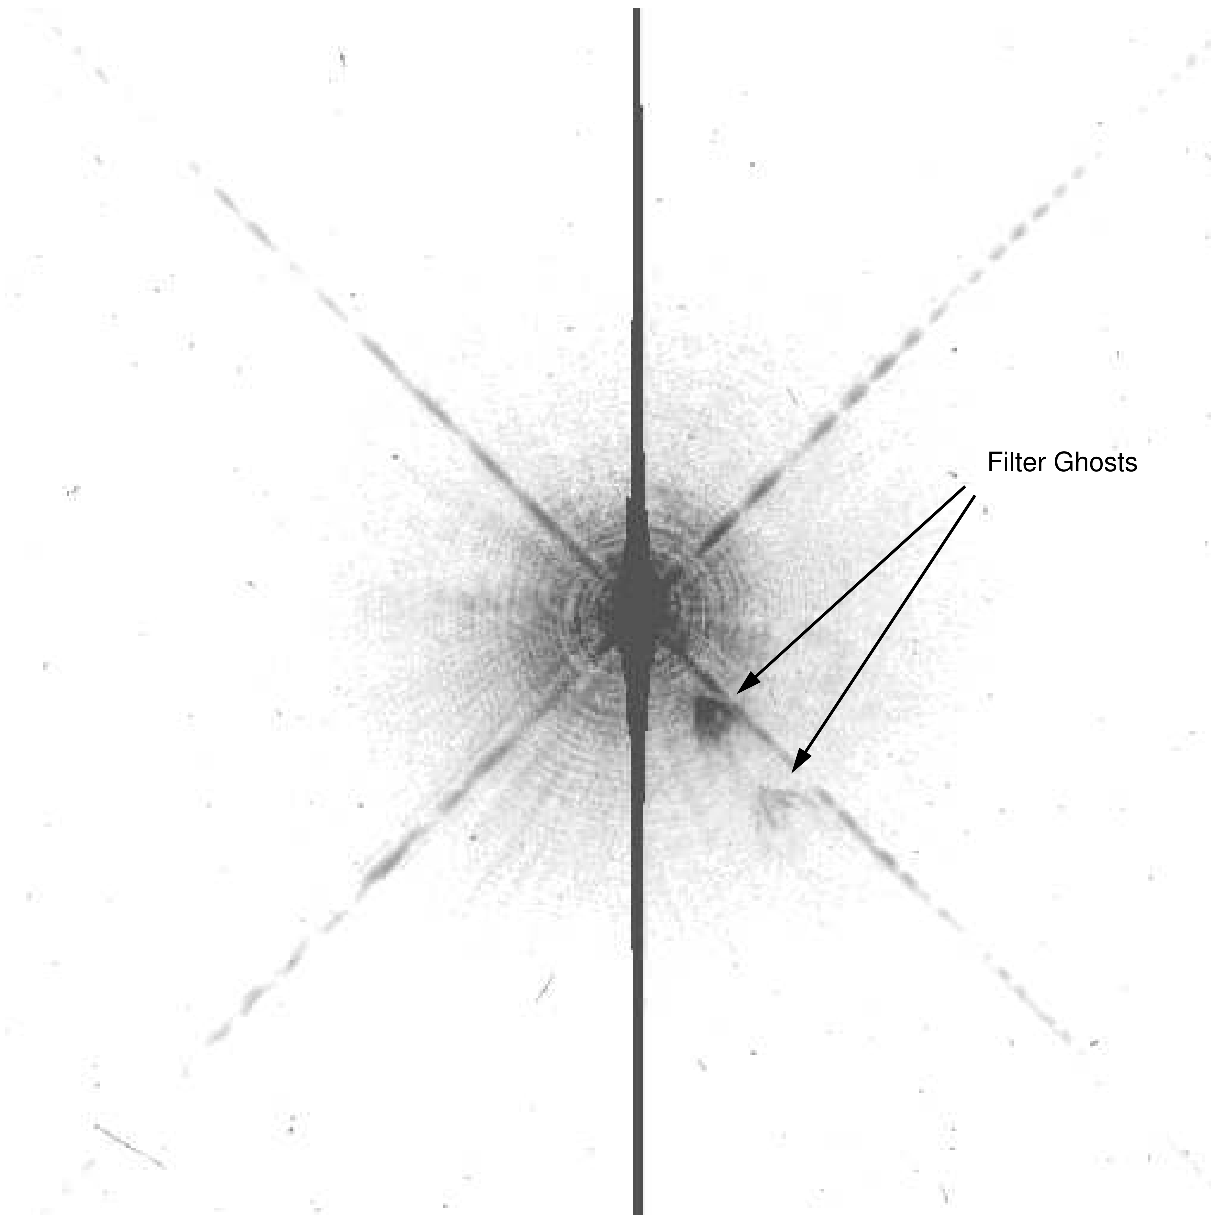

Figure 5.11: Saturated Stellar Image Showing Field Flattener Ghost on WF2.

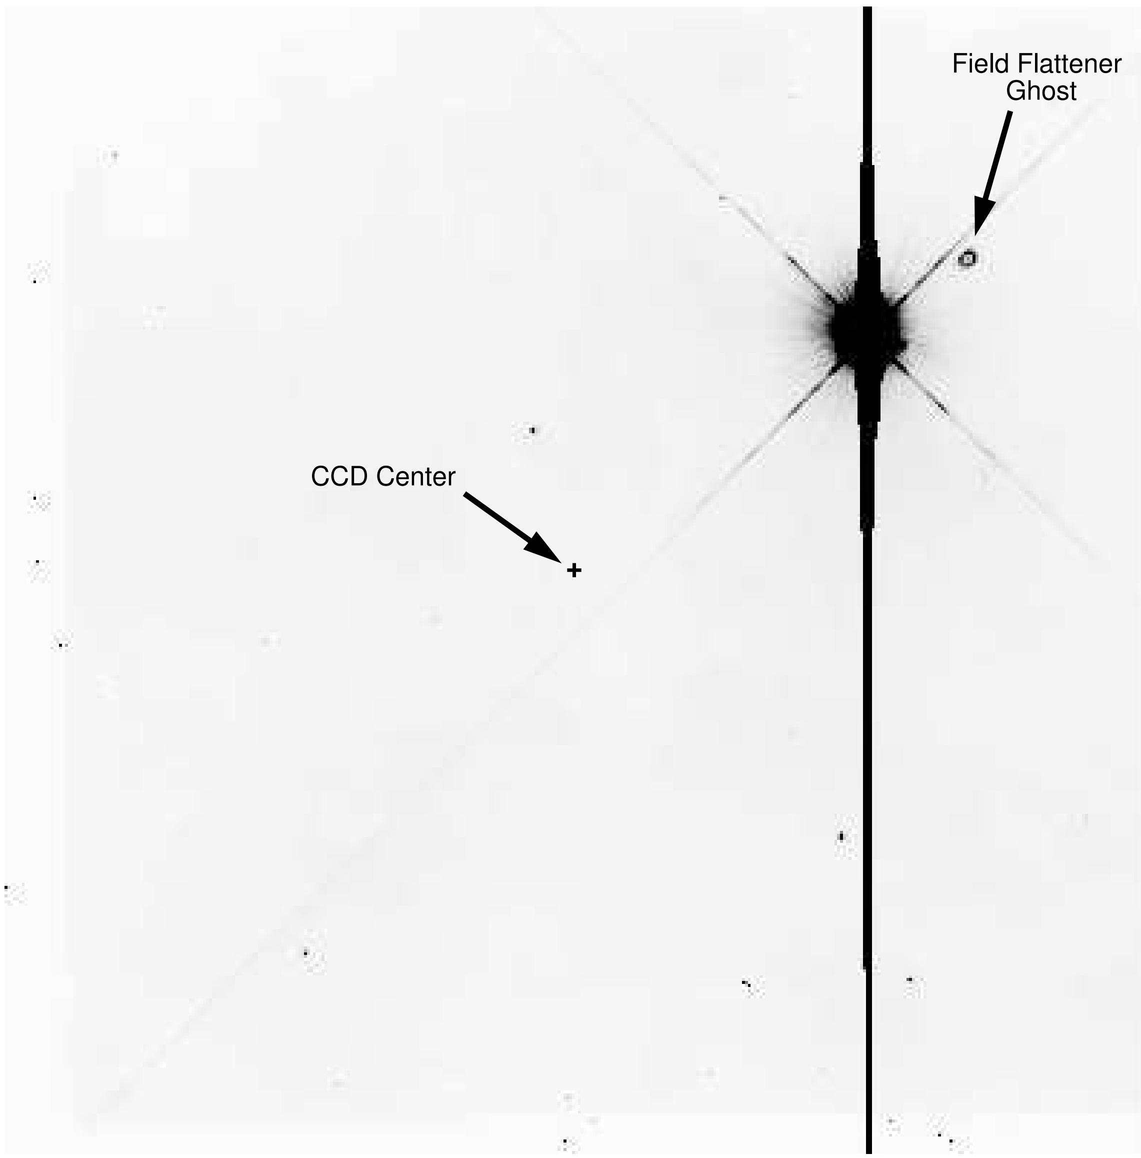

5.10 Optical Distortion

The WFPC2 cameras have significant geometric distortion which not only affects astrometry, but also affects photometry (because the extended sources used to generate flat fields have an induced change in apparent surface brightness). It can be a large effect, with “true” positions differing from observed positions by several pixels in the corners of the cameras. The distortion is wavelength dependent in the ultraviolet, because it is partially caused by the MgF_2 field flattener in front of each CCD. It is sensibly wavelength independent in the visible.

An estimate of the geometric distortion in the WFPC2 cameras was made (Holtzman et al., PASP 107, 156) from a series of 30 F555W exposures in a dense stellar field. Between each of these exposures, the telescope pointing was shifted by 16″ in a 6 x 5 grid along rows and columns. On any given frame, positions of several dozens of stars were measured. Any pair of stars which appeared on three or more of the 30 frames was used to do a cubic distortion solution. The cubic distortion coefficients are of the form:

$$x_{corr} = C_1 + C_2 x + C_3 y + C_4 x^2 + C_5 xy + C_6 y^2 + C_7 x^3 + C_8 x^2 y + C_9 xy^2 + C_{10} y^3$$

$$y_{corr} = D_1 + D_2 x + Dy + D_4 x^2 + D_5 xy + D_6 y^2 + D_7 x^3 + D_8 x^2 y + D_9 xy^2 + D_{10} y^3$$

The coefficients are given in Table 5.5. The input (x,y) in the above equation are offsets from the center of each CCD in pixel units:

$$x = x_{obs} - 400$$

$$y = y_{obs} - 400$$

where (x_{obs}, y_{obs}) are the “observed” pixel positions on each CCD. The corrected values (x_{corr}, y_{corr}) are in a system with the origin near the pyramid apex, and the units are PC1 pixels. Hence PC1 in is quadrant 1, WF2 in quadrant 2, etc. Application of the transformation brings positions of all chips into the orientation of PC1.

The pixel scale can be estimated from the commanded offsets between the frames (relying on the FGS scale and distortion calibrations). It comes out as 0.04554 ± 0.00001 arcseconds pixel^{-1} in the PC, and hence 0.09961, 0.09958, and 0.09964 arcseconds pixel^{-1} in WF2, 3 and 4 respectively. An independent check on an astrometric standard field (M67) yielded 0.04555 arcseconds pixel^{-1} in the PC. These plate scales refer to the scale at the center of the chip in filter F555W. The true scale is lower elsewhere on the chip because of distortion, and there is some wavelength dependence in the scale even for visible wavelengths.

Table 5.5: Cubic Distortion Coefficients.

Coefficient	PC 1	WF2	WF3	WF4
C1	3.54356E+02	-8.12003E+02	-8.07068E+02	7.72904E+02
C2	1.00021 E+00	1.95172E-02	-2.18757E+00	1.37619E-02
C3	9.79758E-04	-2.18805E+00	-8.51284E-03	2.18899E+00
C4	9.84222E-08	-1.51015E-06	-3.21458E-07	1.33326E-06
C5	-6.31327E-09	5.06693E-06	2.77459E-06	-4.00828E-06
C6	-7.19983E-07	1.04063E-06	-2.06874E-06	-5.71282E-07
C7	-3.73922E-08	-4.55094E-10	7.39223E-08	-2.00520E-09
C8	6.65101E-10	7.50249E-08	-1.18833E-09	-7.86673E-08
C9	-3.51470E-08	-1.51652E-09	7.73056E-08	-3.37491E-09
C10	-2.55003E-09	7.34568E-08	-3.72971E-10	-7.81490E-08
D1	3.43646E+02	7.66592E+02	-7.71489E+02	-7.74638E+02
D2	1.00544E-03	2.18607E+00	7.40319E-03	-2.18688E+00
D3	9.99786E-01	2.02006E-02	-2.18554E+00	1.29831E-02
D4	-5.80870E-07	-3.79469E-06	-1.49841E-06	1.30814E-06
D5	-5.07221E-07	-2.99518E-06	2.96650E-06	2.01760E-06
D6	3.41574E-07	1.05789E-07	-2.02756E-07	-8.47656E-07
D7	-1.08091E-09	-7.47857E-08	-1.45720E-09	7.70408E-08
D8	-3.42070E-08	-1.67666E-09	7.72421E-08	-1.68514E-09
D9	1.17819E-09	-7.55302E-08	8.52633E-10	7.75257E-08

The cubic distortion coefficients can be used to derive effective pixel areas as presented in Figure 5.12. Contours are shown at half percent levels. Measurements of total brightness or total counts (as opposed to measurements of surface brightness) should be corrected by multiplying the science image by Figure 5.12. This correction image is also available in the HST data archive as file `f1k1552bu.r9h`.

Figure 5.12: Integrated Photometry Correction Induced by Camera Distortions.