

The Performance and Calibration of WFPC2 on the *Hubble Space Telescope*¹

JON HOLTZMAN,² J. JEFF HESTER,³ STEFANO CASERTANO,⁴ JOHN T. TRAUGER,⁵ ALAN M. WATSON,² GILDA E. BALLESTER,⁶ CHRISTOPHER J. BURROWS,⁴ JOHN T. CLARKE,⁶ DAVID CRISP,⁵ ROBIN W. EVANS,⁵ JOHN S. GALLAGHER III,⁷ RICHARD E. GRIFFITHS,⁸ JOHN G. HOESSEL,⁷ LYNN D. MATTHEWS,⁹ JEREMY R. MOULD,¹⁰ PAUL A. SCOWEN,³ KARL R. STAPELFELDT,⁵ AND JAMES A. WESTPHAL¹¹

Received 1994 July 26; accepted 1994 November 18

ABSTRACT. The WFPC2 was installed in the *Hubble Space Telescope (HST)* in 1993 December. Since then, the instrument has been providing high-quality images. A significant amount of calibration data has been collected to aid in the understanding of the on-orbit performance of the instrument. Generally, the behavior of the camera is similar to its performance during the system-level thermal vacuum test at JPL in 1993 May. Surprises were a significant charge-transfer-efficiency (CTE) problem and a significant growth rate in hot pixels at the original operating temperature of the CCDs ($-76\text{ }^\circ\text{C}$). The operating temperature of the WFPC2 CCDs was changed to $-88\text{ }^\circ\text{C}$ on 1994 April 23, and significant improvements in CTE and hot pixels are seen at this temperature. In this paper we describe the on-orbit performance of the WFPC2. We discuss the optical and thermal history, the instrument throughput and stability, the PSF, the effects of undersampling on photometry, the properties of cosmic rays observed on-orbit, and the geometric distortion in the camera. We present the best techniques for the reduction of WFPC2 data, and describe the construction of calibration products including superbias, superdarks, and flat fields.

1. INTRODUCTION

The Wide Field Planetary Camera 2 (WFPC2) is an instrument for the *HST* that images the central part of the focal plane. It was originally designed to be similar to the WF/PC (hereafter, WF/PC refers to the camera flown in *HST* in 1990, and WFPC2 to the camera installed in 1993 December). Both cameras were built by the Jet Propulsion Lab (JPL). The WF/PC split the central part of the focal plane into four sections with a pyramid-shaped mirror; each of the quadrants

was imaged with an inverted Cassegrain camera onto a TI 800×800 CCD. The WF/PC pyramid could be rotated, allowing the beam to be sent into either a set of $f/13$ or $f/28$ cameras, the Wide Field (WF) or Planetary Cameras (PC). Each set of cameras had its own readout electronics.

With the discovery of spherical aberration, the optics in WFPC2 were modified so that the *HST* pupil would be imaged on the secondaries of the inverted Cassegrains, which were refigured to remove the spherical aberration; this allowed an optical correction without the introduction of any additional surfaces. Actuators were placed on small fold mirrors in three of the four internal cameras and on the main instrument pickoff mirror to allow the beams to be steered accurately onto the corrective optics.

Largely because of budget constraints, WFPC2 only has four internal cameras instead of the eight which were in WF/PC. The decision was made to have one Planetary Camera ($f/28.3$) and three Wide Field Cameras ($f/12.9$). The mechanism for rotating the pyramid was removed. Thus any picture taken with WFPC2 has four quadrants, with one at a higher magnification. Both sets of electronics were retained, giving some redundancy in the signal chains; some new functionality was achieved by configuring the two sets to have different gains to increase the available dynamic range.

In addition to the optical changes, Loral 800×800 CCDs were used in place of the original TI devices. This change was motivated largely because of concerns about the quality of the small remaining stock of TI devices. The Loral CCDs do not have the problem of quantum-efficiency hysteresis (QEH) that was present in the TI devices, and consequently, UV-flooding is not needed. In addition, the Loral CCDs have lower readout noise ($\sim 5e^-$) and uniform response across the field. The two gains provided by the electronics are $7e^-/\text{DN}$ (electronics bay 4) and $14e^-/\text{DN}$ (electronics bay 3). In the $14e^-/\text{DN}$ channel, the effective readout noise mea-

¹Based on observations with the NASA/ESA *Hubble Space Telescope*, obtained at the Space Telescope Science Institute, which is operated by AURA, Inc., under Contract No. NAS 5-26555.

²Lowell Observatory, Mars Hill Road, Flagstaff, Arizona 86001. Electronic mail: holtz@lowell.edu, alan@lowell.edu

³Department of Physics and Astronomy, Arizona State University, Tyler Mall, Tempe, Arizona 85287. Electronic mail: jjh@cygloop.la.asu.edu, scowen@tycho.la.asu.edu

⁴Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218. Electronic mail: stefano@stsci.edu, burrows@stsci.edu

⁵Jet Propulsion Laboratory, 4800 Oak Grove Drive, Mail Stop 179-225, Pasadena, California 91109. Electronic mail: jtt@wfp2-mail.jpl.nasa.gov, dc@crispy.jpl.nasa.gov, rwe@wfp2-mail.jpl.nasa.gov, krs@wfp2-mail.jpl.nasa.gov

⁶Department of Atmospheric, Oceanic, and Space Sciences, University of Michigan, 2455 Hayward, Ann Arbor, Michigan 48109. Electronic mail: gilda@sunshine.sprl.umich.edu, clarke@sunshine.sprl.umich.edu

⁷Department of Astronomy, University of Wisconsin–Madison, 475 N. Charter St., Madison, Wisconsin 53706. Electronic mail: jsg@jayg.astro.wisc.edu, hoessel@jth.astro.wisc.edu

⁸Department of Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, Maryland 21218. Electronic mail: griffiths@mds.pha.jhu.edu

⁹Astronomy Program, State University of New York at Stony Brook, Stony Brook, New York 11794. Electronic mail: matthews@sbastl.ess.sunysb.edu

¹⁰Mount Stromlo and Siding Springs Observatories, Australian National University, Private Bag, Weston Creek Post Office, ACT 2611, Australia. Electronic mail: jrm@mso.anu.edu.au

¹¹Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, California 91125. Electronic mail: jaw@sol1.gps.caltech.edu

sured from frames is $\sim 7.5e^-$ because of digitization effects. Either gain can be selected by the observer.

A more detailed description of the instrument can be found in the WFPC2 Instrument Handbook (1994) and the WFPC2 Science Calibration Report (1993).

This paper provides a description of the on-orbit status of the WFPC2 and those issues that affect the reduction of WFPC2 data. This paper builds on the understanding of the instrument gained during the system level thermal vacuum test ("TV") at JPL in 1993 May. Section 2 gives a brief summary of the thermal history, optical history, and status of proposals. Section 3 describes the WFPC2 calibration pipeline at STScI, including the products used in the pipeline to date and the time scale and prognosis for improvement of these products. Section 4 describes the WFPC2 point-spread function (PSF). Section 5 describes the photometric performance of the instrument, including a discussion of charge transfer properties, a description of the time-dependence of throughput and contamination, a presentation of observed throughput and a comparison of the observed throughput with expectations, a discussion of photometric calibration for the broad band filters, and data relating to subpixel QE variations. Section 6 provides a characterization of the cosmic-ray properties observed on-orbit. Section 7 discusses the geometric distortions in the cameras.

This document represents our understanding of the WFPC2 as of fall 1994. This understanding continues to evolve, sometimes rapidly, and any of the discussions or results presented herein are subject to revision.

2. INSTRUMENT STATUS

2.1 Thermal History of CCDs

The WFPC2 CCDs were designed to operate between -70 and -90 °C. At these temperatures, they are among the coolest components on the *HST* and as such, the CCD windows are a site for the condensation of contaminants. Contaminants absorb most strongly in the UV. For this reason, the buildup of contaminants was carefully monitored during the initial cooldown by regular UV observations of a bright star. No evidence of contamination was seen. The CCDs were first cooled to the nominal operating temperature of -76 °C around December 26, and remained at this temperature until February 22.

Contaminants do build up slowly, and over a time scale of several weeks degradation of the UV throughput can be seen (see Sec. 5.2). To restore the UV throughput, the CCDs are warmed approximately once per month to about 20 °C for about 6 h. This procedure seems to completely remove the contaminants. The dates of these decontamination procedures through fall 1994 are listed in Table 1.

The *HST* had a safing event on April 11 and again during the first week of July. During these events the WFPC2 CCDs remained at their operating temperature and there was essentially no impact on the instrument.

During March and April, we discovered a significant charge transfer efficiency (CTE) problem in the WFPC2 CCDs (discussed in Sec. 5.1). Because lab data indicated that the problem might be significantly reduced at a colder CCD

operating temperature, the WFPC2 CCDs were cooled from -76 to -88 °C after the decontamination procedure which occurred on 1994 April 23. As discussed below, the new operating temperature significantly improves the instrument CTE, and it seems likely that the CCDs will be kept at the new temperature. Consequently, 1994 April 23 marks a critical date in the history of WFPC2. Some of the calibration products and the photometric zeropoints are temperature dependent, so it is important to use the correct calibration products during data reduction.

2.2 Optical History

Initial images were taken between 18 and 25 December 1993. Based on these, several small changes were made to the telescope focus and the tilts of the WFPC2 pickoff mirror and articulated fold mirrors (AFMs). After these moves, the image quality was judged to be sufficiently good to allow early release observations (EROs) and regular science observations to proceed. This optical alignment was used from 1993 December 28 until 1994 March 4.

To determine final adjustments to the optics, some images were taken far from focus and phase retrieval was used to determine the small remaining aberrations (Burrows and Krist 1994). As a result of this analysis, small adjustments were made to the WFPC2 optics on 1994 March 4. These changes are expected to have only a very minor effect on the calibration of the instrument.

The *HST* focus continues to change because of desorption of the telescope structure, but the rate of this change is now quite small (Hasan and Burrows 1994). Small focus adjustments are made infrequently to compensate for this. A small (5μ) motion of the secondary was made on 1994 June 29.

Data from before and after the servicing mission indicate that the *HST* focus has short time scale variations corresponding to secondary motions of order $\sim 5-10 \mu$ that are presumably related to the thermal state of the telescope. This behavior manifests itself as a change in the PSF on orbital time scales, and has been described as "breathing."

A more detailed description of *HST* and WFPC2 optical parameters is given in Burrows and Krist (1994). A discussion of the WFPC2 PSF is presented in Sec. 4.

2.3 Scheduling History

After the first set of test images, ERO images were taken during the week of 1993 December 26. These included images of M100, η Carinae, two QSO host galaxies, and 30 Doradus. Servicing Mission Orbital Verification (SMOV) calibration observations designed by the WFPC2 IDT (Investigation Definition Team) were taken during early 1994. These included photometric monitors, initial photometric calibration observations, internal monitors with biases, darks, and internal flats, external flat-field observations, and some optical characterization programs. Around mid-March, the SMOV programs finished and the STScI Cycle 4 calibration and monitor observations started. Some effort was made to ensure continuity between the SMOV and Cycle 4 calibration observations, although in some cases targets were changed. Calibration data taken during Cycle 4 are being

TABLE 1
Decontamination Dates

Date
22 February 1994
24 March 1994
23 April 1994
23 May 1994
13 June 1994
10 July 1994
28 July 1994
27 August 1994
24 September 1994
20 October 1994

jointly analyzed by the WFPC2 IDT and the STScI. All calibration data are public and available to users.

3. CALIBRATION STATUS

The WFPC2 calibration pipeline at the STScI consists of the following steps:

- (1) Correction of small analog-to-digital conversion (ADC) errors;
- (2) Subtraction of a bias level;
- (3) Subtraction of a superbias frame;
- (4) Subtraction of a superdark frame;
- (5) Correction for shutter shading effects;
- (6) Application of a flat field.

In this section, we discuss the calibration products (biases, darks, flat fields, etc.) used in the STScI pipeline and also discuss additions to the standard calibration that may be of interest to some observers. Generally, it takes some time to update calibration products in the STScI pipeline, so it is often the case that data automatically processed by the pipeline are not reduced with the most up-to-date products. All calibration products are eventually made available and users who wish to reduce their own data can obtain the most appropriate products from STScI.

The raw data supplied by STScI consist of two frames: the 800×800 data frames (DOH files) and the 800×14 engineering frames (XOH) files. The engineering frames consist of encoded engineering data (the first two columns) plus an overscan region of 12 columns which contain bias data read through the signal chain. STScI also distributes pipeline-reduced frames (COH files) and additional spacecraft and instrument engineering data (SHH files).

3.1 ADC Error Correction

The WFPC2 ADCs produce small errors in the conversion from electrons to DN. We believe that this is related to the use of switching power supplies which may generate noise that influences the ADC conversion in a systematic way. The

ADC corrections for WFPC2 are much smaller than the corresponding corrections for WF/PC; the largest correction is ~0.1 percent.

The first step in WFPC2 image reduction is to correct for these systematic ADC errors by replacing the raw output DN values with the expected mean input values. Currently, corrections derived from TV data are being used and there is a separate correction for each gain state. Monitoring of the ADC problem throughout TV showed that the correction was stable. It is hoped that internal ramp filter flats taken on-orbit will provide a check of the ADC correction. If any significant changes are seen, a new ADC correction will be derived.

The ADC correction should be applied to both the image data and the overscan data.

3.2 Bias Level Subtraction

The next step is to subtract a bias level from the data. As discussed in the WFPC2 TV report, the first few overscan columns are affected by the level of exposure on the chip. Consequently, for accurate bias subtraction, the first few overscan columns should be ignored. The effect of using all 12 columns is that a slightly elevated overscan subtraction is subtracted. This can lead to significant errors in the inferred background and to flat-fielding errors. These effects are most important for faint objects.

There is evidence for small differences in bias level between the even columns and the odd columns. It appears that the difference between odd and even columns in bias frames varies with time and can fall in one of two separate states. The differences are typically less than 0.1 DN; however, in one state in WF4, the difference is about 0.4 DN. Because the effect can reach this amplitude, we believe that a rigorous reduction should use separate odd and even bias values. We average columns 9, 11, and 13 (one origin) in the engineering frame to get the even bias level in the data frame and use columns 10, 12, and 14 to get the odd bias level (note the mismatch between even and odd between the engineering and data frames).

Until 1994 March, all 12 overscan columns were used in the STScI pipeline; since early March only the last 6 overscan columns have been used. The STScI pipeline originally used a single bias value for all columns, but now uses a separate odd and even bias subtraction.

3.3 Superbias Subtraction

Some structure remains in frames of zero exposure time after a uniform (or a separate odd-even) bias subtraction. This structure includes real bias structure and also some dark current. Dark current is present because the interval between reset and read is always at least one minute. To remove this structure, we subtract superbias frames.

The superbias frames originally in the pipeline were derived from TV data from medians of a large number of individual bias frames after subtraction of mean bias levels. Bias frames are being taken as part of the WFPC2 internal-monitor program, and superbias frames have been derived from these data. A separate superbias is made for each gain state (because different electronics are used) and for each chip (because each

chip has a unique bias and dark signature). The superbias also depends on CCD temperature, so separate superbias frames have been constructed for -76°C and for -88°C . In superbias construction, we subtract a separate odd and even column bias value from each frame as described above, so the resultant superbias should be free of odd-even column effects.

There is some evidence for long-term and intermittent variations in the bias structure at sub-DN levels. Some frames have horizontal bias jumps which can run for several tens of rows, others have vertical streaks, and there is also occasional evidence for a horizontal odd-even effect. All of these effects occur at the sub-DN level and are not important for most applications. We have not devised any methods for removing these low-level variations, but users should be cautioned that they may exist in reduced data, and we recommend that data be carefully inspected for their presence.

3.4 Superdark Subtraction

Every CCD pixel has an individual dark rate. To remove the dark current, we construct superdark frames from many long dark frames. In TV, we confirmed that the observed superdarks for the two gain states differed only by a factor of two, the ratio of the gains. Consequently, the superdark in the pipeline for both channels is derived from the low gain channel in which digitization effects are less severe (i.e., the superdark for the bay 3, gain=14 channel is the superdark for the bay 4, gain=7 channel divided by two).

The growth rate of hot pixels on-orbit, especially at the -76°C operating temperature, is significantly greater than expected from ground tests. Furthermore, some of the hot pixel disappear (“anneal”) when the chips are warmed during the monthly decontaminations.

The change in the population of hot pixels requires the construction of new superdark frames on a fairly frequent basis. The STScI current calibration plan takes ~ 5 darks every week and these darks are available to observers. Every few weeks, the STScI pipeline is populated with a new superdark generated from darks taken over the preceding few weeks. Consequently, the superdarks in the pipeline at the time data are reduced may not be appropriate. Also, because these superdarks are made from only a relatively small number of frames, the statistics do not allow a very accurate measurement of the dark rate. The noise in the determination of the dark rate for pixels which vary cannot be reduced, but a more accurate dark rate can be derived for pixels which do not vary.

We have constructed superdarks by combining all of the darks taken at -76°C and, separately, the first 20 darks taken at -88°C . In this process, we rejected outlying bright pixels to remove cosmic rays. These superdarks accurately measure the dark rate for pixels which did not vary. For pixels which varied in dark rate, the superdark values are representative of the state in which the pixels spent the most time.

In Fig. 1 we show the dark rate for “typical” pixels in each chip from the on-orbit superdarks; data from the -76°C and -88°C superdarks are shown with solid and dotted

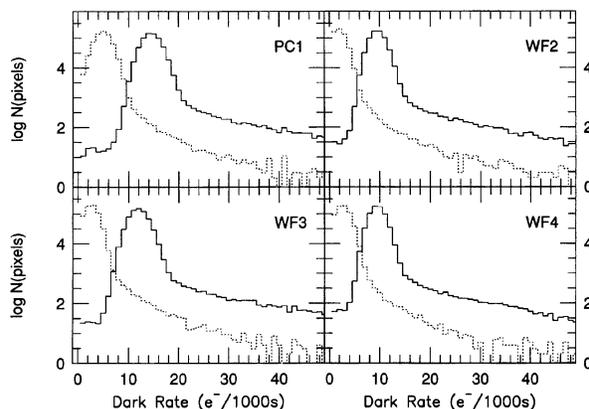


FIG. 1—The dark rate for pixels in each chip from the on-orbit superdarks. Data from the -76°C and the -88°C frames are shown with solid and dotted lines.

lines. Note that the dark rate is shown in $e^-/1000s$, not in DN. The goal for WFPC2 was $10e^-/1000s$. At -76°C , the mean rates for PC1 and WF3 are $\sim 14e^-/1000s$; for WF2 and WF4, the rates are $\sim 10e^-/1000s$. These numbers are consistent with those measured in TV. At -88°C , the dark rate is significantly lower and easily meets the goal.

To investigate the time dependence of darks, we constructed “local darks” from dark frames taken over periods of roughly one week. The top panels of Fig. 2 show the dark rates for each local dark from two months of operation at -76°C (top left) and the first month at -88°C (top right). These plots show the cumulative number of pixels with rates larger than the dark rate on the abscissa. The lowest dark rates are not shown in Fig. 2 because they cannot be determined accurately from the small number of frames used to make the local darks (because of readout noise and low-level cosmic rays).

There is clearly a great variation in the number of hot pixels with time, especially at -76°C . At any given time, the observed number of hot pixels N with rates above a certain value x is fairly well-approximated by a power law:

$$N(\text{rate} > x) = N_0 x^{-\alpha},$$

where $\alpha \sim 1$ and the proportionality constant N_0 varies with time. This representation fails for the hottest pixels.

In the top panels of Fig. 2, data from the first set of darks taken immediately after the first decontaminations at each temperature are shown with bold solid lines. A significant fraction of the hot pixels disappear during a decontamination during which the CCDs are held at $\sim 20^\circ\text{C}$ for ~ 6 h. It seems that the hottest pixels disappear fairly completely. At -76°C , many pixels with dark rates of a few tens of $e^-/1000s$ remain. The cumulative number of pixels with rates greater than 35 (top line), 70, 140, 280, and 560 (bottom line) $e^-/1000s$ are shown as a function of time in the middle panel of Fig. 2. Decontaminations dates are marked with vertical dotted lines. The epoch when the CCDs were cooled to -88°C (94.31) is apparent from the significant drop in the number of hot pixels. The number of pixels with rates greater than $35e^-/1000s$ is dominated by those with

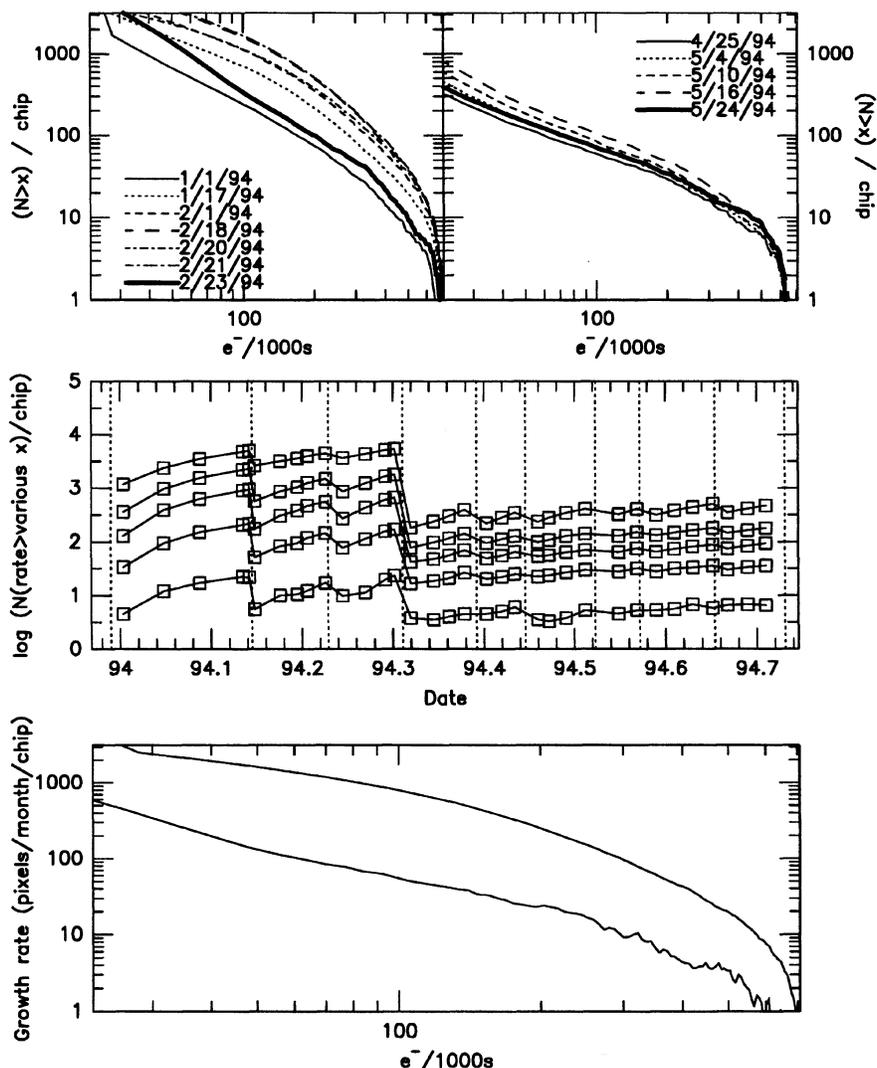


FIG. 2—Data for hot pixels in local darks. The top panels show the total number of pixels with rates greater than the values specified on the abscissa. Various lines show data from different epochs during the first decontamination cycle at -76°C (left panel) and -88°C (right panel). The bold lines show the first set of data after the first decontaminations at each temperature. The middle panel plots the number of hot pixels as a function of time for pixels with rates greater than 35 (top line), 70, 140, 280, and 560 (bottom line) $e^{-}/1000\text{ s}$. The dotted lines show times of decontamination procedures; note that the chip temperature was lowered to -88°C after the decontamination procedure around 94.31. The bottom panels show the growth rate of hot pixels (number of pixels/month) inferred from data during the first decontamination cycle at each temperature; the higher curve is for -76°C , the lower for -88°C .

rates just a few times the goal of $10e^{-}/1000\text{ s}$. At -76°C , the total number of these pixels appears to be growing with time even with the regular decontaminations. At -88°C , however, there is currently no strong evidence for the long term accumulation of hot pixels.

There is a clear break in the behavior of the hot pixels at the point that the CCD temperature was changed to -88°C . At this lower temperature, the number of hot pixels is down by roughly an order of magnitude. The typical growth rate of hot pixels at the two operating temperatures is shown in the bottom panel of Fig. 2. At -88°C , several hundred hot pixels appear each month; at -76°C , the rate was several thousand pixels per month. Although hot pixels are created and

anneal at -88°C , there are ≤ 1000 hot pixels per chip at any time, corresponding to less than 0.1 percent of the chip area. Consequently, hot pixels are not nearly as much of a problem at the current operating temperature as when the CCDs were at -76°C .

For the best calibration, especially at -76°C , darks taken at a time near the observation are required. In principle, such darks can be used to correct hot pixels, although the rate of appearance of hot pixels is sufficiently large that a fair number of hot pixels appear even over the course of a day. However, these darks, by necessity, will be constructed from only a few frames and so Poisson noise will limit the accuracy of the derived dark rates.

Also, in order for a dark subtraction to be useful, the dark rate for any pixel must be constant in time. We have inspected a subset of pixels which became hot between 1994 January 1 and February 23 to determine whether this is true. We find that a good number ($\sim 50\%$ – 75%) of the pixels have variations which are consistent with Poisson statistics, but a significant fraction have larger variations. The larger variations come in several forms: some pixels go to one rate for a given time, then switch to another rate (either higher or lower); some pixels go hot for a while then return to a “normal” state; and other pixels just have random (either a single or multiple) deviations. The latter may be artifacts of the process by which local darks are constructed. Because of dark rate variations, we caution that even the use of a local dark can give erroneous dark subtractions, with errors above those expected from Poisson statistics alone. In investigations that require a good understanding of the noise, we suggest that all hot pixels should be ignored, rather than corrected.

Our preferred method for dark subtraction is to use well-determined super darks derived from a large number of frames for all pixels except those which show a significantly different dark rate in the local darks closest to the date of the observations. We either ignore these pixels or estimate their dark rates from the local darks, depending on the science to be done. There are some subtleties in this process and we recommend that observers contact the STScI instrument scientists or the authors before attempting this themselves.

The most straightforward approach, and one which probably works well at -88°C is to simply scale and subtract superdark frames. Since WFPC2 is only read out on 60 s boundaries and since the superbias includes some dark current, the time used to scale the superdark is

$$t_{\text{sd}} = s + 60 \times \text{ifix}[(t + 16.4)/60],$$

where t is the requested open shutter time, s is 60 for exposures taken with the serial clocks off (default for long exposures) and 0 for exposures taken with serial clocks on, and the *ifix* function represents integer truncation.

3.5 Shutter Shading Correction

The shutter causes small variations in the exposure time across the field of view. A shutter shading correction was determined during TV. The maximum peak-to-peak error across the entire WFPC2 field of view is about 5 millisecc. This leads to at most a 5 percent error in a 0.11 s exposure, and the effect is negligible for exposures of more than a few seconds. Nonetheless, the TV data were used to construct shutter shading correction frames, and these were used to implement a shutter shading correction in the pipeline after roughly 1994 March 1. There are two shutter blades, and consequently a separate correction depending on which blade is used. It is unlikely that on-orbit data can provide as accurate a measurement of this effect as the extensive TV data, so the TV data will be used in the pipeline indefinitely. The corrections derived in TV appear to work well on on-orbit observations of the earth, which are taken with very short

exposure times. In addition, the on-orbit engineering data on shutter timing is essentially the same as that obtained during TV.

3.6 Flat Fields

3.6.1 Strategy

One of the most significant problems with the calibration of the original WFPC was the determination of flat fields (Faber et al. 1991). The only external flat-field source available was the bright earth. While various algorithms were devised which successfully removed the “streaks” from earth flats, the bright earth was far too bright to allow flats to be obtained through the broadband “workhorse” photometric filters without crossing them first with a neutral density filter. Unfortunately, inclusion of the ND filter (as well as “reciprocity-like” effects in the CCDs) significantly compromised the success of this effort.

The selection of the Loral CCDs for WFPC2 enabled us to take a very different approach to obtaining flat fields for the new camera. The Loral CCDs are very stable devices. They do not need to be UV flooded, and lab testing indicates that the characteristics of their flat fields do not change appreciably over time or across thermal cycling. This stability was verified during TV, and has since been verified by comparing internal calibration channel flats taken on-orbit with similar flats taken during TV.

We made use of the stability of the detectors to devise a flat-field strategy in which the flat-field patterns of all filters were well characterized during TV, and tied together in a way that would allow them to be transferred to orbit using a minimum of on-orbit data. Specifically, significant effort was made to ensure that a uniform and *wavelength independent* source of illumination was available during TV. Analysis of TV data showed that this goal was accomplished for visible wavelengths ($\geq 3500 \text{ \AA}$). This source was used to build up very high signal-to-noise flat fields through all filters. These flats contained all information about the variations in CCD response and spatial structure in filter transmission, and were expected to differ from on-orbit flats principally by the ratio of the *HST*+WFPC2 illumination pattern to the TV illumination pattern.

3.6.2 Pipeline Flats before 1994 Mid-March

Prelaunch estimates of flats were constructed from the TV flats using an empirical model of the TV illumination pattern and a theoretical model of the *HST* illumination pattern. Specifically, the TV illumination pattern was “measured” by assuming that the CCDs themselves had little low-frequency structure in their sensitivity. The TV flats were then multiplied by the ratio of an optical model of the *HST* illumination pattern to the TV illumination pattern. These flats were installed in the pipeline prior to launch and remained there until 1994 mid-March. Comparison of new flats (next section) with these original flats shows that these flats were generally good to a few percent over most of the field, with a maximum deviation of about 10 percent. Some small errors at the level of a few percent occurred in chip-to-chip normalization of the original pipeline flats.

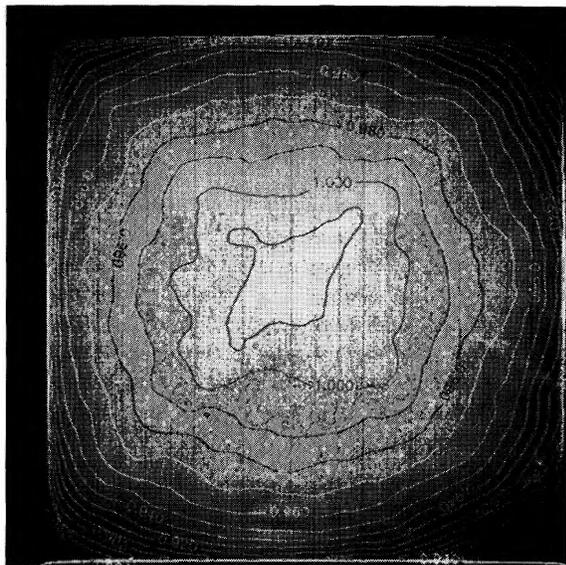


FIG. 3—The *F555W* flat for WF2. The greyscale ranges from 0.8 to 1.2 of the mean intensity.

3.6.3 On-Orbit Visible Flats

Improvement on the pre-launch flats involved the measurement of the *HST*+WFPC2 illumination pattern. This was accomplished using earth flats taken in a number of narrow-band filters. These earth flats were flattened using the pre-launch flats and corrected for shutter shading effects. Observations that showed evidence for significant streaks (from features on the earth) were rejected and the remainder were combined using a simple clipped mean after normalization.

The *HST* illumination pattern has been derived independently using *F375N*, *F502N*, and *F953N*, and by simultaneous reduction of data from all three filters. The rms residuals in the comparison of the *HST* illumination pattern derived from each filter with the pattern derived from all filters simultaneously are between 0.3% and 0.6%. This is consistent with the expected uncertainties in these flats on the basis of the level of streaking in individual frames and the volume of data available. Peak-to-peak differences among these flats were generally <2%. This supports the assumption that both the *HST* and TV illumination patterns are wavelength independent. On this basis, the average *HST* illumination pattern was smoothed and applied to all of the TV flats. These flats were normalized using the average over the region (200:600,200:600) in WF3 and then inverted (so frames should be multiplied by the pipeline flat fields).

A picture of one of the *F555W* flats is shown in Fig. 3 to show the appearance of a typical flat. Flats for WFPC2 are remarkably uniform and do not vary strongly with wavelength. The small features come from dust and from spots in the lumogen coating. One can see the “X” at the center of the field which occurs where the obscuration from the WFPC2 internal secondary supports line up with those of the *HST*, causing a region of slightly higher throughput.

3.6.4 Caveats

There are a number of caveats that apply to the current flats. The first is that the flat fields were constructed using TV and on-orbit data taken at a CCD temperature of -76°C , whereas the current operating temperature is -88°C . Fortunately, while there may be variations in *absolute* QE of a few percent between -76 and -88°C , the *spatial structure* in the flats is constant between these two temperatures at a level well less than 1%. (Differences between TV flats at -76 and -88°C were consistent with photon statistics.)

A second caveat involves structure in the flats associated with features on the pyramid, which is near a focal plane. Such features include dust and the edges of the pyramid. Because of slight changes in alignment, these structures shifted between TV and orbit. Care was taken when deriving the *HST* illumination pattern to shift the locations of these features. However, this was done in a wavelength-independent fashion, whereas the features themselves are somewhat wavelength dependent due both to diffraction and refraction by the CCD field flatteners. As with any camera, care should be taken when interpreting data near the location of in-focus obscurations. The difference in alignment between TV and orbit also results in a slight “ripple” pattern in the ratio of TV to on-orbit flats arising from slight variations in the reflectivity of the pyramid. This pattern was correctly removed in the construction of the on-orbit flats.

The current generation of on-orbit flats are normalized to unity in the region (200:600,200:600) of WF3. Flats for all chips are cross normalized to give the same number of counts (DN)/area in each chip for a given source. We believe that the current cross-chip normalizations are good to a couple of percent; they may be adjusted slightly in the future. Observers are cautioned that photometric zero-points are tied to a particular set of flats.

A final caveat involves the interaction between filter wedge and the alignment of the WFPC2, TV stimulus, and *HST* pupils. Since the *HST* and TV stimulus secondary supports (spiders) are on opposite sides of the entrance pupil, differences in the wedges of the filters cause the spider shadows to move in different ways relative to the location of the WFPC2 secondary supports. This results in a slight filter dependent difference in the location of the “X” in the flat field where the obscurations line up. Analysis of TV data in which the alignment was intentionally changed indicates that this effect should be <0.5% for most WFPC2 filters.

3.6.5 Ultraviolet Flats

The situation is less ideal for the UV flats than for visible light flats because neither the TV illumination source nor the internal calibration channel have wavelength independent illumination patterns. Furthermore, it is difficult, if not impossible, to get far-UV flats of the bright earth. A further complication is that the WFPC2 CCDs derive their UV and blue sensitivity from a lumogen phosphor coating which converts UV and blue photons to visible photons.

To construct UV flats, we combined three different components: high spatial frequency, pixel-to-pixel variations for each UV filter, a low-frequency lumogen response, and the

HST+WFPC2 illumination pattern. The lumogen response and *HST*+WFPC2 illumination patterns were assumed to be independent of wavelength. We compared TV flats at 4100 Å to those around 5500 Å to derive the lumogen response. The high spatial frequency components were derived from internal calibration channel flats by dividing out the lumogen pattern and then removing low spatial frequencies separately for each filter. The *HST*+WFPC2 illumination pattern for these flats was assumed to be the same as that derived for the visible flats.

Unfortunately we have no data to allow us to reliably distinguish between variations in the calibration channel UV illumination pattern and true spatial variations in filter transmission. Analysis of TV data indicates that while this is probably not a serious problem at wavelengths longer than about 2500 Å, there are uncertainties of several percent for shorter wavelengths. Furthermore, there is evidence for significant structure in the red leaks of the conventional UV filters. Comparison of *F218W* earth flats with *F218W* TV flats shows a smooth 30% ramp between the outer corners of WF2 and WF4. *F218W* frames taken during TV with different color lamps indicate that this gradient is almost certainly due to structure in the red leak of the filter. Significant improvement in UV flats and better characterization of spatial variations in red leak await additional data such as observations in which both red and blue stars are stepped across the field of the WFPC2.

3.6.6 Independent Tests of Flat Field Quality

Several groups have attempted to use the large data sets collected by the Medium Deep Survey team (MDS) or the GTO parallel program to derive “sky flats” for *F606W* and *F814W*. The MDS team reports that the agreement between the sky flats and the pipeline flats is very good, with no obvious problems at the $\pm 2\%$ level over most of the field of the WFPC2 (Casertano 1994, private communication). The greatest exception to this is seen within about 30 pixels of the edges of the CCDs, where differences of up to 5% between the two flats are seen. Interestingly, structure is seen along the CCD edges when directly comparing flats from other sources such as TV flats or calibration channel flats. These are probably due to the combined effects of slight alignment changes, the aberrated PSF present in the *HST* focal plane, and the sharp edges of the pyramid facets. It may be possible, at least for *F606W* and *F814W*, to improve the pipeline flats using the sky flat data.

Slight differences ($\leq 1\%$) are also seen between the MDS and pipeline flats in the vicinity of the “X” where the *HST* and WFPC2 secondary supports line up. As noted above, small errors are expected here in the pipeline flats due to the differences in filter wedge from filter to filter coupled with differences in the behavior of the TV stimulus and *HST* pupils. On the other hand, there is also reason to believe that these differences could result from scattered light in the sky flats. Analysis of deep frames taken as part of the GTO parallel program show the spider pattern changing with time in a way that seems to correlate with the position of the bright earth (Groth 1994, private communication). In the case of earth flats a strong source of extraneous light is certainly

present, but it is generally no brighter than the source illuminating the proper field of the telescope. In the sky flats, however, very low levels of extraneous light can be a significant fraction of the total source. Thus, it is unclear whether the pipeline flats or the sky flats provide the best characterization of the structure of the “X” for the *F606W* and *F814W* filters.

While slight differences exist between different efforts, the general conclusion reached from comparison of the independently derived pipeline and sky flats is that our knowledge of the WFPC2 flat field is good. It seems unlikely, with the exception of very localized features, that uncertainty in flat fields currently represents a significant limitation in the photometric capability of the instrument. Efforts continue both to monitor the flats using the internal calibration source and to extend the earth and sky flat data sets.

3.6.7 The Effect of Distortion on Photometry and Flat Fields

The WFPC2 cameras have significant geometric distortion as a result of their optical design. This distortion causes the effective pixel area to vary across the field, which has an interesting effect on integrated photometry of flat-fielded images. The flat fields created on-orbit by their nature correct for the pixel area effect; the low-frequency structure in the flats includes a component which comes from effective pixel area. Consequently, after flattening, all pixels are normalized to the same area. Photometrically, this preserves surface brightness. However, the geometric correction which is included in the flat fields renders total brightnesses incorrect. For accurate integrated photometry, it is necessary to multiply by the effective pixel areas. This correction is small over most of the field but can reach $\sim 4\%$ in the corners of each CCD. The measured distortion and pixel area maps are presented in Sec. 7.

3.6.8 Edge Effects in Flat Fields

The WFPC2 focal plane is split into four quadrants by the pyramid mirror. However, this mirror is located before the optics which correct for spherical aberration. The beam at the pyramid mirror has significant spatial extent, so consequently, light from a given location on the sky can fall into more than one of the WFPC2 quadrants. Each quadrant is then imaged with aberrated optics onto a CCD. The net result is that the edges of the pyramid shadow visible in each of the WFPC2 CCDs has the shape of a spherically aberrated PSF collapsed along one dimension; a plot of this shape (from a PSF model with spherical aberration) is shown in Fig. 4. The details of this edge depend sensitively on fold mirror tilts, filter wedge, chromatic aberration, and any other factor which affects the details of the internal alignment of the instrument. Because the alignment was adjusted, the location of the edge moved between the TV and on-orbit flats.

The current flats attempt to give a uniform illumination all the way to the chip boundaries. A wavelength independent edge correction was included to correct the TV flats for the movement of the edge. This average treatment is neither exact for each individual filter nor is it perfectly constant in time. Furthermore, since there is very little light at the

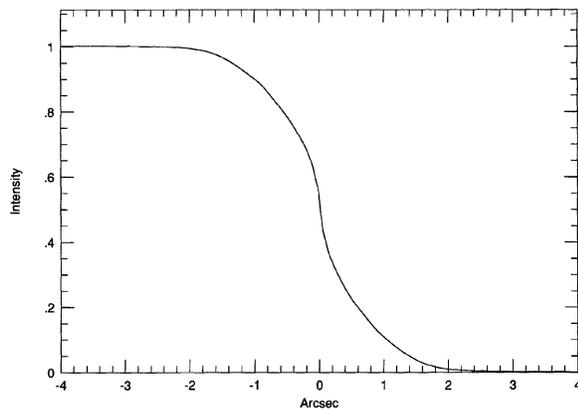


FIG. 4—The shape of the edge of a flat field in WFPC2, determined from a model of a spherically aberrated PSF collapsed along one dimension.

boundaries, the flats are less accurate in this region. Figure 4 should give observers some idea of the correction that the flat field makes as a function of distance from the edge. Users are cautioned that within an arcsecond or so of the edge, the flats may be in error by a few percent.

4. THE WFPC2 POINT-SPREAD FUNCTION

The primary goal of WFPC2 was to correct the spherical aberration of the *HST*. This goal was accomplished and the

WFPC2 optics appear to deliver nearly diffraction-limited images.

Observed PSFs vary with wavelength, field position, and time. Variations with field position arise because of a variable pupil function as well as from small field-dependent aberrations. Variations with time occur from variable spacecraft jitter during exposures and from focus changes, which can occur on short time scales (the “breathing” mentioned earlier).

4.1 Encircled Energy Curves

We have accumulated a set of “mean” encircled energy curves from eight separate observations in eight different filters for two of the cameras, PC1 and WF3. All of these data are for stars near the center of the chips and for short exposures. Even still, there is some variation between different observations, arising from jitter, focus changes, photon statistics, pixel centering, errors in background determination, and throughput variations. The WFPC2 PSFs, like all others, extend essentially to infinity. The far-field PSF is dominated by scattering from small-scale structure on the mirrors. An estimate of the magnitude of this scattering has been made from observations of bright stars (Krist and Burrows 1994), and has been found to be significantly larger in WFPC2 than in WF/PC. As it is essentially impossible to measure the far-field PSF from observations of unsaturated stars, we have chosen to normalize the encircled energy curves presented here to the light within a $1''$ radius circular aperture.

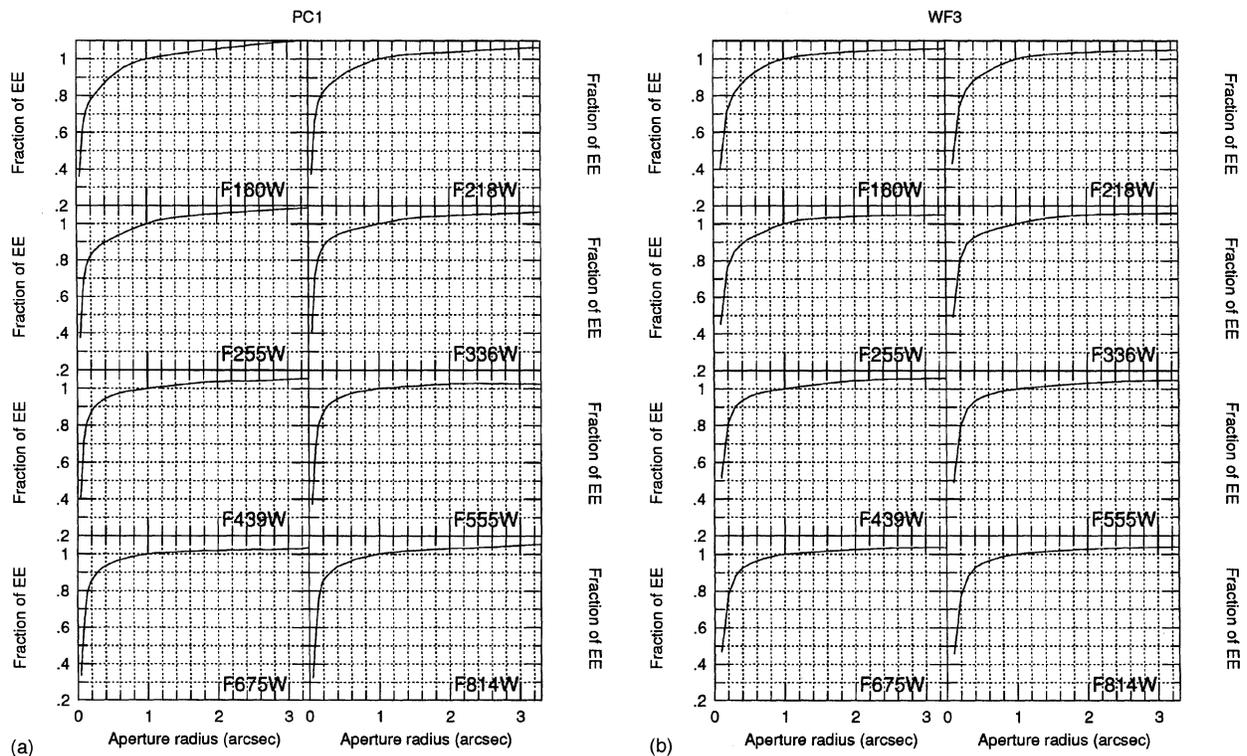


FIG. 5—(a) Typical encircled energy curves for PC1 for several filters. For each filter, the mean of eight different observations is shown. (b) Typical encircled energy curves for WF3 for several filters. For each filter, the mean of eight different observations is shown.

The mean curves are plotted in Figs. 5(a) and 5(b) (for the two different chips) and are tabulated, with errors, in Tables 2(a) and 2(b); the errors in these tables are the standard deviation from the set of eight observations. For the visible filters (*F439W*, *F555W*, *F675W*, and *F814W*), the curves are normalized to the mean value at 1". For the shorter wavelengths, throughput variations are significant (see Sec. 5.2), so each curve was normalized to its own value at 1"; consequently, the errors at this radius are zero by definition. The sky value for these curves was taken from an annulus defined by 6" and 8" radii. For the smallest apertures in each chip, the values are sensitive to the technique used for including fractional pixels; for the data here, we used a technique from DAOPHOT (Stetson 1987), in which the circular aperture is approximated with an irregular polygon.

Tables 2(a) and 2(b) show that there can be several percent variation in the PSFs at small radii, even near the centers of the chips. We have estimated additional scatter which arises from field dependent PSF variations using PSF models. We find that field variations are likely to introduce an additional 1%–2% scatter for small radii (less than a few pixels). Within a few arcseconds of the edge, the PSF variations are more severe, as discussed below.

Observers are cautioned that blind application of the values listed in Tables 2(a) and 2(b) may introduce systematic errors, especially at small radii. They are provided here to give information on the general character of the PSF, and can be used for approximate aperture corrections for images in which it is impossible to determine such corrections directly from the image (e.g., very crowded fields).

4.2 CCD Effects on the WFPC2 PSF

At small radii, the encircled energies are significantly higher in the PC than in the WFs. We believe that this is caused by a "pixel scattering" effect in the CCDs which causes a fraction of the light which falls entirely in one pixel to leak into neighboring pixels. Since a higher fraction of light falls in the central pixel in the WFs, this effect causes encircled energies in the WFs at small radii to be degraded compared to the PC. We measure that roughly 20%–30% of the light incident on one pixel leaks into adjacent pixels.

4.3 Edge Effects on the WFPC2 PSF

Within 3 arcsec of the pyramid edges, light from a star will fall into more than one of the WFPC2 CCDs because the beam is spherically aberrated at the pyramid. The WFPC2 optics correct the spherically aberrated beam, but near the edge only part of the pupil will be filled in each camera. Consequently, the PSF will be distorted near the edges of the chips. For example, at the point at which half the light falls into each chip, the pupil is semicircular; diffraction from such an aperture produces an elongated PSF. For this reason, and because of uncertainties in the flats, photometry near the edge of the pyramid may be uncertain.

5. PHOTOMETRIC STABILITY AND CALIBRATION

5.1 Charge Transfer Effects

During January, we observed a field of stars in ω Centauri (ω Cen) at different pointings such that a given star appeared at many (~ 30) different locations in the different exposures. This test was originally designed to test the flat-fielding accuracy and to understand the field variations of the PSF in the cameras. When we reduced the data, we found significantly more scatter in the multiple observations than we had expected. Inspection of the residuals showed a clear trend in the sense that a star was measured to be fainter when it appeared at higher row numbers. There is no evidence for any monotonic structure along columns in the flat fields. Also, the effect is present along columns in all four chips, even though each chip is rotated by 90° in the focal plane.

We believe that this phenomenon is caused by charge transfer problems within the CCDs. After discovery on-orbit, this effect was duplicated in the lab using similar CCDs. It is believed that this effect arises from traps in the silicon which keep electrons from being transferred during readout. The trends that have been seen can be explained if the accessibility of these traps depends on the total number of electrons being passed through each pixel. Note that this is not a "classical" deferred charge effect; some of the traps can be filled only if there is a large number of electrons passing through the pixel, so the effect cannot be entirely removed by applying a preflash. However, we expect that the effect is significantly reduced in the presence of background light on the CCDs. It is worth noting that this effect was *not* present in the prelaunch tests of CTE because the source used for CTE testing produced charge over a large enough fraction of the CCD to quickly saturate the traps and allow essentially lossless transfer.

The upper panel of Fig. 6 shows the residuals (mean-observed magnitude) of the ω Cen photometry at -76°C as a function of row number; the mean magnitudes were calculated from all of the different observations of the star, and only stars with at least 10 observations are included. In addition to the dithered data, we also have data obtained for photometric calibration in a field which has been observed from the ground (see Sec. 5.6). In this field, we have some estimate of "truth" from the ground observations. The left panels of Fig. 7 plots the residuals in three different colors (ground-flight magnitudes) as a function of row number, including data from all four chips.

These data show a 10%–15% gradient in photometry along the columns of each chip. Further analysis shows that the effect has a different amplitude for stars of different brightnesses; brighter stars (i.e., stars with several thousand DN) have only a $\sim 5\%$ gradient, while fainter stars have a larger gradient. There is some difference between the WFs and the PC, and also between different filters, as might be expected for a charge transfer effect since that would presumably depend on the number of pixels over which light from a star is spread.

Lab data suggested that the CTE effect would be significantly reduced by operating the CCDs at a colder temperature. This motivated the decision to change the WFPC2 CCD set points to -88°C on 1994 April 23. A shortened dither test was performed at this temperature, and results from this are shown in the lower panel of Fig. 6. In addition, the data from the calibration field at -88°C are shown in the right panels of Fig. 7. The effect is significantly reduced at the colder operating temperature, although it is not gone entirely; at -88°C , the effect is $\sim 3\%$ – 4% peak-to-peak across the entire chip. The flight hardware does not allow the chips to be operated any colder than -88°C .

We are working on a semiempirical model that can be used to correct for this effect. At the present time, there is only a limited amount of data with which we can constrain such a model. We are currently obtaining additional and improved lab data. One of the largest uncertainties at this point is the behavior of the CCDs if there is some background present on the chip. All the ω Cen exposures were short and contain very little background, but many science exposures are long and contain significant background. One set of relevant observations is in the galaxy M101, which was observed as part of the H_0 Key Project at both CCD operating temperatures at the same orientation. Photometry of this field, which has a typical background level of 200 – $300 e^-$,

shows no differences between the two temperatures as a function of row (Han 1994, private communication), implying that CTE effects with this background are similar and probably negligible at both temperatures. At high background levels there is no apparent effect judging from the flat-field exposures.

The effect is probably complicated, with the loss in any given pixel dependent on the values on other pixels in the same column, as well as on the value in the pixel in question. Until a more sophisticated algorithm for correction of the effect is available, observers may wish to consider the following simple correction. For observations with little background (≤ 20 – $30 e^-$), apply a 4% ramp from the bottom to the top of each CCD, leaving the values of pixels at low rows as they are read out, and increasing the apparent brightness of those at high rows by 4%. For observations with intermediate background levels ($\sim 30 e^-$ to $\sim 250 e^-$), apply a 2% ramp. For higher background levels, apply no correction. Note that even without any correction, the CTE effect will only introduce $\sim 1\%$ rms scatter at the current operating temperature for applications in which objects are spread over the entire field. For observations of an object at a single location, there could be up to a $\sim 4\%$ error in the absolute photometry if the CTE effect is not taken into consideration.

TABLE 2(a)
Approximate EE Curves for PC1¹

Rad ²	F160W	F218W	F255W	F336W	F439W	F555W	F675W	F814W
1	0.31±0.03	0.38±0.06	0.37±0.04	0.41±0.04	0.41±0.02	0.37±0.07	0.34±0.04	0.32±0.01
2	0.58±0.02	0.65±0.03	0.66±0.03	0.71±0.02	0.72±0.02	0.67±0.09	0.61±0.06	0.58±0.01
3	0.68±0.02	0.75±0.02	0.77±0.01	0.81±0.01	0.82±0.01	0.80±0.05	0.79±0.05	0.77±0.01
4	0.73±0.01	0.80±0.01	0.82±0.01	0.86±0.00	0.86±0.01	0.85±0.02	0.85±0.02	0.85±0.00
5	0.76±0.01	0.83±0.01	0.84±0.01	0.88±0.00	0.89±0.01	0.89±0.01	0.87±0.01	0.87±0.00
6	0.79±0.01	0.85±0.01	0.86±0.01	0.90±0.00	0.91±0.01	0.91±0.01	0.90±0.01	0.89±0.00
7	0.82±0.01	0.86±0.01	0.88±0.01	0.92±0.00	0.93±0.01	0.92±0.01	0.92±0.01	0.91±0.00
8	0.85±0.01	0.88±0.01	0.89±0.01	0.93±0.00	0.94±0.01	0.94±0.01	0.93±0.01	0.92±0.00
9	0.87±0.01	0.89±0.01	0.90±0.01	0.94±0.00	0.95±0.01	0.95±0.01	0.94±0.01	0.93±0.00
10	0.89±0.01	0.90±0.01	0.91±0.00	0.95±0.00	0.95±0.01	0.95±0.01	0.95±0.01	0.94±0.00
11	0.91±0.01	0.92±0.01	0.92±0.01	0.95±0.00	0.96±0.01	0.96±0.01	0.95±0.01	0.95±0.00
12	0.92±0.01	0.93±0.01	0.93±0.01	0.96±0.00	0.97±0.01	0.97±0.01	0.96±0.01	0.95±0.00
13	0.94±0.00	0.94±0.01	0.94±0.01	0.96±0.00	0.97±0.01	0.97±0.01	0.97±0.01	0.96±0.00
15	0.96±0.00	0.95±0.00	0.95±0.00	0.97±0.00	0.98±0.01	0.98±0.01	0.98±0.01	0.97±0.00
20	0.99±0.00	0.99±0.00	0.99±0.00	0.99±0.00	1.00±0.01	1.00±0.01	1.00±0.01	0.99±0.01
25	1.01±0.00	1.01±0.00	1.02±0.00	1.01±0.00	1.01±0.02	1.01±0.01	1.01±0.02	1.01±0.01
30	1.02±0.00	1.02±0.01	1.03±0.00	1.03±0.00	1.02±0.02	1.01±0.01	1.01±0.02	1.02±0.01
35	1.03±0.01	1.03±0.01	1.04±0.00	1.04±0.00	1.03±0.02	1.02±0.01	1.02±0.02	1.02±0.01
40	1.04±0.00	1.04±0.01	1.05±0.01	1.04±0.01	1.03±0.02	1.02±0.01	1.02±0.03	1.03±0.01
45	1.05±0.01	1.04±0.01	1.06±0.01	1.05±0.01	1.04±0.03	1.03±0.01	1.02±0.03	1.03±0.02
50	1.05±0.01	1.04±0.01	1.06±0.01	1.05±0.01	1.04±0.03	1.03±0.01	1.02±0.04	1.03±0.02
55	1.06±0.01	1.05±0.02	1.07±0.01	1.05±0.01	1.04±0.03	1.03±0.01	1.03±0.05	1.04±0.03
60	1.07±0.01	1.05±0.02	1.08±0.02	1.06±0.01	1.04±0.04	1.03±0.01	1.02±0.05	1.04±0.03
70	1.08±0.02	1.06±0.03	1.09±0.02	1.06±0.03	1.05±0.05	1.03±0.02	1.03±0.05	1.05±0.04
80	1.09±0.02	1.07±0.03	1.10±0.02	1.07±0.05	1.06±0.06	1.02±0.02	1.04±0.06	1.06±0.05
90	1.10±0.02	1.08±0.04	1.10±0.02	1.08±0.06	1.07±0.06	1.01±0.03	1.04±0.08	1.06±0.07
100	1.11±0.03	1.08±0.04	1.12±0.03	1.08±0.06	1.07±0.06	1.01±0.04	1.05±0.09	1.07±0.08
110	1.12±0.03	1.08±0.05	1.12±0.03	1.08±0.06	1.07±0.07	1.00±0.04	1.05±0.10	1.08±0.10
120	1.12±0.04	1.09±0.06	1.13±0.03	1.08±0.06	1.08±0.08	1.00±0.05	1.06±0.12	1.08±0.11
130	1.13±0.04	1.09±0.07	1.15±0.04	1.09±0.06	1.08±0.08	0.99±0.06	1.06±0.13	1.09±0.11

¹The encircled energies here were determined by averaging eight separate observations of a single flux standard. Accurate aperture corrections are a function of both time and location on the chip, and as such, these numbers only give an approximation to “true” aperture corrections. The errors are the standard deviation of the eight measurements.

²The aperture radius is given in PC pixels. The radius in arcsec can be determined by multiplying by the scale of $0.0455''/\text{pixel}$.

TABLE 2(b)
Approximate EE curves for WF3¹

Rad ²	F160W	F218W	F255W	F336W	F439W	F555W	F675W	F814W
1	0.30±0.05	0.43±0.06	0.45±0.06	0.49±0.06	0.51±0.05	0.49±0.09	0.47±0.07	0.46±0.06
2	0.63±0.04	0.74±0.04	0.76±0.04	0.80±0.03	0.81±0.02	0.80±0.05	0.78±0.05	0.77±0.05
3	0.77±0.01	0.83±0.01	0.85±0.01	0.89±0.01	0.90±0.01	0.90±0.02	0.88±0.02	0.88±0.01
4	0.84±0.01	0.88±0.01	0.89±0.00	0.92±0.00	0.94±0.01	0.93±0.01	0.93±0.01	0.93±0.01
5	0.89±0.01	0.91±0.00	0.92±0.00	0.95±0.00	0.96±0.01	0.95±0.01	0.95±0.01	0.95±0.01
6	0.92±0.01	0.93±0.00	0.94±0.00	0.96±0.00	0.97±0.01	0.97±0.01	0.96±0.01	0.96±0.01
7	0.95±0.00	0.95±0.00	0.95±0.00	0.97±0.00	0.98±0.01	0.98±0.01	0.98±0.01	0.98±0.01
8	0.97±0.00	0.97±0.00	0.97±0.00	0.98±0.00	0.99±0.01	0.99±0.01	0.99±0.01	0.99±0.01
9	0.99±0.00	0.99±0.00	0.99±0.00	0.99±0.00	0.99±0.01	1.00±0.01	0.99±0.01	0.99±0.01
10	1.00±0.00	1.00±0.00	1.00±0.00	1.00±0.00	1.00±0.01	1.00±0.01	1.00±0.01	1.00±0.01
11	1.01±0.00	1.01±0.00	1.01±0.00	1.01±0.00	1.01±0.01	1.00±0.01	1.00±0.01	1.00±0.01
12	1.01±0.00	1.02±0.00	1.02±0.00	1.02±0.00	1.01±0.01	1.01±0.01	1.01±0.02	1.01±0.01
13	1.02±0.00	1.02±0.00	1.02±0.00	1.02±0.00	1.02±0.01	1.01±0.01	1.01±0.02	1.01±0.01
14	1.02±0.00	1.02±0.00	1.03±0.00	1.03±0.00	1.02±0.01	1.01±0.01	1.01±0.02	1.01±0.01
15	1.03±0.00	1.03±0.00	1.03±0.00	1.04±0.00	1.03±0.01	1.02±0.01	1.01±0.02	1.02±0.01
16	1.03±0.00	1.03±0.00	1.03±0.00	1.04±0.00	1.03±0.01	1.02±0.01	1.02±0.02	1.02±0.01
17	1.04±0.00	1.03±0.00	1.04±0.01	1.04±0.00	1.03±0.01	1.02±0.01	1.02±0.02	1.02±0.01
18	1.04±0.00	1.03±0.00	1.04±0.01	1.05±0.00	1.04±0.01	1.03±0.01	1.02±0.02	1.02±0.01
19	1.04±0.00	1.04±0.01	1.04±0.01	1.05±0.00	1.04±0.01	1.03±0.01	1.02±0.02	1.03±0.01
20	1.04±0.00	1.04±0.01	1.04±0.01	1.05±0.00	1.04±0.01	1.03±0.01	1.03±0.02	1.03±0.01
21	1.05±0.01	1.04±0.01	1.04±0.01	1.05±0.00	1.05±0.01	1.03±0.01	1.03±0.02	1.03±0.01
22	1.05±0.01	1.04±0.01	1.04±0.01	1.05±0.00	1.05±0.02	1.04±0.01	1.03±0.02	1.03±0.01
25	1.05±0.01	1.04±0.01	1.05±0.01	1.05±0.01	1.05±0.02	1.04±0.01	1.03±0.02	1.03±0.01
30	1.06±0.01	1.05±0.01	1.05±0.01	1.06±0.01	1.06±0.02	1.05±0.01	1.04±0.02	1.04±0.01
35	1.06±0.01	1.05±0.01	1.05±0.02	1.06±0.01	1.06±0.02	1.05±0.01	1.04±0.02	1.04±0.02
40	1.06±0.01	1.05±0.02	1.05±0.02	1.06±0.01	1.06±0.02	1.05±0.01	1.04±0.02	1.04±0.02
45	1.07±0.02	1.05±0.02	1.05±0.02	1.06±0.01	1.07±0.02	1.05±0.01	1.04±0.02	1.04±0.02
50	1.07±0.02	1.06±0.02	1.06±0.03	1.06±0.01	1.07±0.03	1.06±0.01	1.04±0.02	1.04±0.02
55	1.08±0.02	1.06±0.02	1.06±0.03	1.07±0.01	1.07±0.03	1.06±0.02	1.05±0.03	1.04±0.02
60	1.08±0.03	1.07±0.03	1.06±0.03	1.07±0.01	1.07±0.03	1.06±0.02	1.05±0.03	1.04±0.03

¹The encircled energies here were determined by averaging eight separate observations of a single flux standard. Accurate aperture corrections are a function of both time and location on the chip, and as such, these numbers only give an approximation to “true” aperture corrections. The errors are the standard deviation of the eight measurements.

²The aperture radius is given in WF pixels. The radius in arcsec can be determined by multiplying by the scale of 0.0996"/pixel.

5.2 Stability of Throughput Inferred from External Sources

The photometric performance of WFPC2 is significantly more stable than that of WF/PC because of significant improvements in contamination control. Despite these improvements, however, it was always expected that there would be some degradation of throughput in the far-UV arising from the accumulation of contaminants on the cold CCD windows. Because of this, we periodically warm the CCDs to remove the accumulated contaminants.

To monitor this degradation, exposures of a UV-bright flux standard (AGK +81 266) through the F170W filter were taken every 3 days through most of the SMOV period (proposal 4721). The measured count rate for this star in each CCD is shown in Fig. 8(a). Dotted lines in this plot mark dates of decontaminations when the CCDs were warmed. This plot shows an initial phase where the throughput was fairly constant for a few days, an increase around date 93.99 which corresponds to the time when the CCDs were warmed for COSTAR deployment, then a series of “sawteeth” where the throughput declines for a few weeks, then is restored during a decontamination. The rate of degradation is variable between the chips, with a maximum of about 10 percent per month for the first two months. In sub-

sequent months, the rate of degradation appears to have lessened, suggesting that there were some contaminants from the servicing mission which have escaped from the instrument. This rate of contamination is roughly two orders of magnitude lower than that observed in WF/PC in the UV (WF/PC OVSV report).

Figure 8(a) shows that after decontaminations, the UV throughput is completely restored. Consequently, we infer that the contaminant(s) causing the degradation in UV throughput are removed during decontaminations.

The final points for WF3 and WF4 in Fig. 8(a) (filled symbols) are from data taken after the chip temperatures were changed to -88°C ; these observations were saturated in PC1 and WF2. The throughput at -88°C appears higher than it ever had been at -76°C . This effect is most likely caused by an improvement in charge transfer effects.

After cooling to -88°C , we transitioned to a different monitoring proposal and a different flux standard (GRW +70 5824). The data points which have been obtained for this star are shown in Fig. 8(b). At -88°C , a higher rate of contamination is seen ($\sim 1\%$ loss of throughput per day in F170W). This is not surprising as contaminants are expected to accumulate more rapidly on a colder surface. We will continue to monitor the throughput degradation at -88°C through sev-

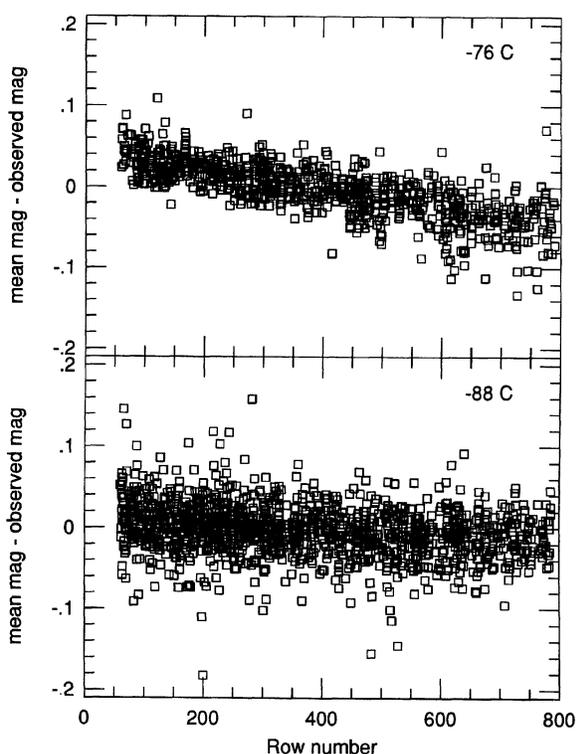


FIG. 6—Mean-observed magnitudes of stars in different pictures in the ω Cen dither test as a function of row number. The top panel shows data at -76°C ; only stars which appeared in at least 10 (of 30) different exposures are included. The bottom panel shows results from the abbreviated dither test at -88°C ; only stars which appear in at least 6 (of 9) different exposures are included.

eral filters. Based on these data, it should be possible to provide an approximate correction to photometry at any given epoch. The frequency of decontaminations will probably remain at about once per month.

Another set of proposals (4763 and 5563) monitor throughput in two of the four chips (PC1 and WF3) through additional UV filters as well as through a few visible filters; these two proposals had an overlap period but used different flux standards (4763 used BD +75 320, 5563 uses GRW +70 5824). Results for several filters are shown in Figs. 9(a) and 9(b) for the two different proposals; in Fig. 9(b), data from both CCD temperatures are shown, with -88°C as filled symbols. The left axes show the absolute count rate as measured through an aperture 1 arcsec in radius while the right axes show the relative throughput with a normalization to the maximum observed count rate. Note that plots for different filters have different scales in order to show effects in all filters. The formal error bars on these points are smaller than the point size; true errors are probably larger because of the presence of low-level remnant cosmic rays (visible cosmic rays were removed manually before measuring). Figures 9(a) and 9(b) provide the data which can be used to estimate the rate of contamination as a function of wavelength.

These data show a clear degradation at wavelengths from 1200 to 3500 \AA with an increasing rate of degradation at

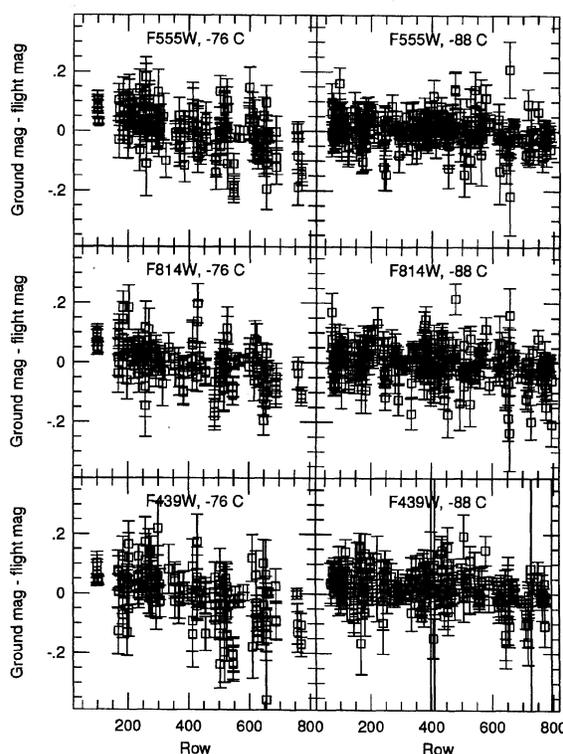


FIG. 7—Ground-observed on-orbit magnitudes of ω Cen stars as a function of row number at -76°C (left panels) and -88°C (right panels). Data are shown for three different filters: $F555W$ (top), $F814W$ (middle), and $F439W$ (bottom). Color terms and zero points were solved for in the transformation from the on-orbit to the ground data; consequently, the mean value on-orbit is equal to the ground-based mean in each passband.

shorter wavelengths. There is possibly evidence for a small amount of degradation at 5500 \AA . The data at long wavelengths are fairly sparse and noisy, and even if degradation is present, it is only at the $\sim 1\%$ per month level. The throughput at all wavelengths appears to be completely restored by a decontamination.

Since the data are only of a single star, it is not possible to tell whether the degradation has spatial structure (but see Sec. 5.3).

5.3 Stability of Throughput Inferred from Internal Flats

Images taken with the internal calibration channel have also been used to monitor contamination of the WFPC2 CCDs. While the lamp stability is not adequate to allow the absolute QE to be monitored, variations in the structure of the internal flats can be seen. We were alerted to these variations by the OSS staff, who reported a complex of *bright* “worm-like” features in $F255W$ calibration channel flats taken on 1994 February 14. These features are typically a few hundred pixels long and a few dozen pixels across. They are most pronounced along the right edge of WF2, but a few features are present in all CCDs.

Subsequent analysis showed that these features were in the same locations as *dark* features seen in $F170W$ flats

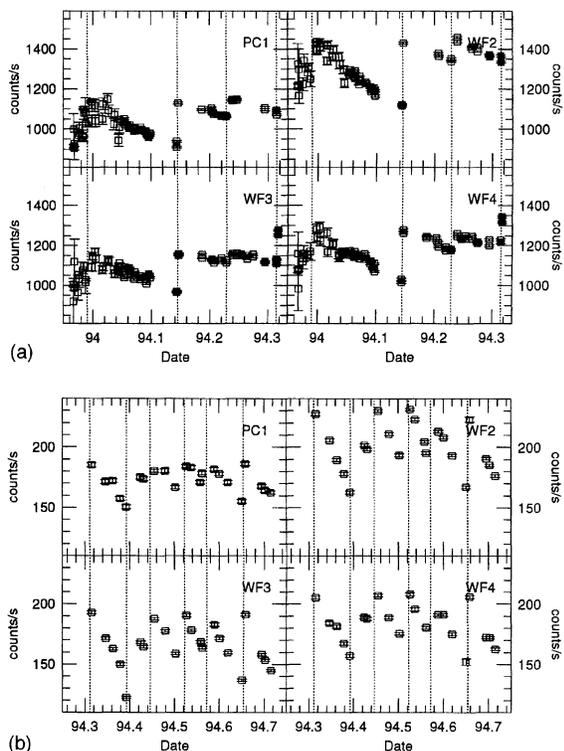


FIG. 8—(a) UV throughput through the $F170W$ filter as a function of time for each of the four CCDs, from observations of the star AGK +81 266. Dates of decontamination procedures are shown with dotted lines. The solid points (last set on WF3 and WF4) were taken at -88°C ; all others were at -76°C . All measurements were made with a circular aperture of radius 3 arcsec. (b) UV throughput through the $F170W$ filter as a function of time for each of the four CCDs, from observations of the star GRW +70 5824. Dates of decontamination procedures are shown with dotted lines. All points are from data at -88°C . All measurements were made with a circular aperture of radius 3 arcsec.

during TV. These dark $F170W$ features *disappeared* over the period of time on-orbit that the bright $F255W$ features appeared. Following decontamination the worms “reset” to the TV state. The most reasonable model of these features is that they trace the location of some permanent contaminant on the CCD windows that absorbs at $F170W$ (to explain the dark TV features), but on which other contaminants do *not* stick very well (to explain the evening out of the $F170W$ features and the appearance of bright $F255W$ features). The “color” changes in these features indicate that the absorption spectrum of the worm material differs significantly from the absorption spectrum of the time-dependent contaminants.

At $F255W$, the variations in response between the worms and the rest of the chip can grow to be about 2% after one month at -76°C . After this amount of time, the worms can be faintly seen at longer wavelengths, but their amplitude is significantly lower; less than 1% at optical wavelengths. Consequently, while the worms represent an interesting effect and provide useful information for tracking contamination, they are probably of little concern for the analysis of science data. It is possible that observers will be able to see the presence of worms in optical data taken at the end of the monthly decontamination cycles, however, the maximum

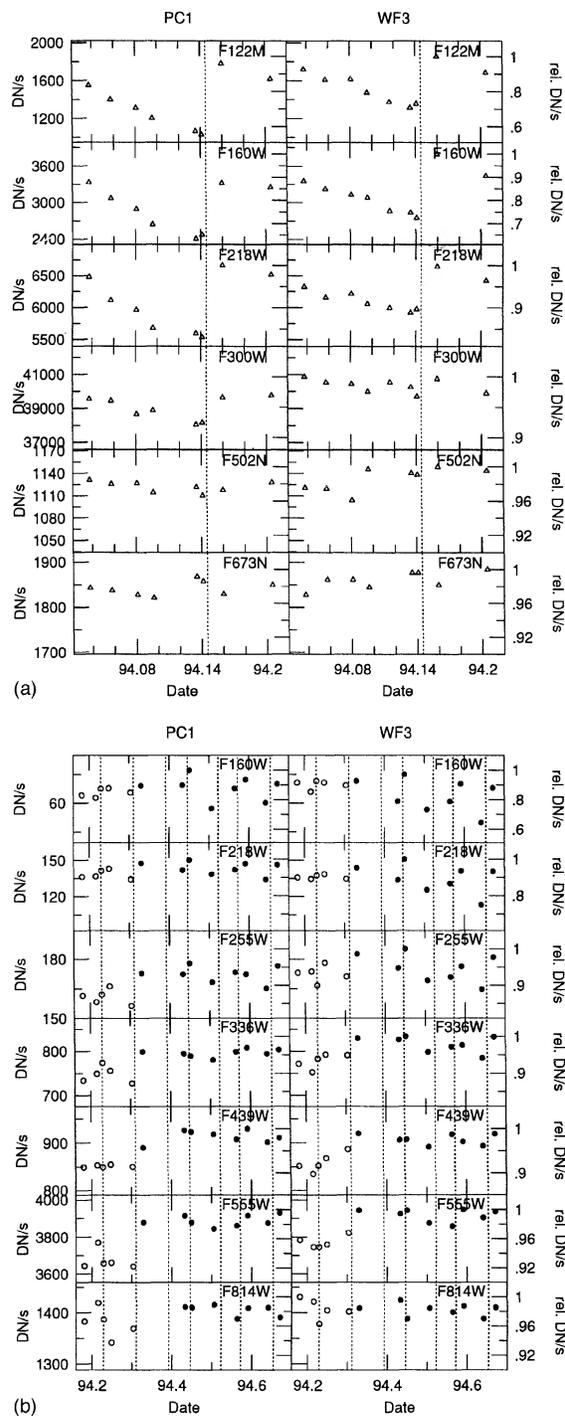


FIG. 9—(a) Throughput through a variety of filters as a function of time for the star BD +75 320. All data are with a CCD temperature of -76°C . Dates of decontamination procedures are shown with dotted lines. The left axis shows absolute count rates; the right axis show relative rates, with the normalization set to unity at the maximum brightness observed. (b) Throughput through a variety of filters as a function of time for the star GRW +70 3824. Data with a CCD temperature of -76°C are shown with open points; data at -88°C are shown with filled points. Dates of decontamination procedures are shown with dotted lines. The left axis shows absolute count rates; the right axis show relative rates, with the normalization set to unity at the maximum brightness observed.

TABLE 3
Percent Change in Internal Flat Level, Referenced to TV Data

Filter	Date	PC1	WF2	WF3	WF4
F555W	19/1/94	-0.12	0.06	0	0.00
	24/1/94	-0.16	0.02	0	-0.02
F439W	14/2/94	-0.23	-0.23	0	-0.01
	19/1	-0.15	0.12	0	0.04
F255W	24/1	-0.16	0.03	0	0.00
	14/2	-0.17	-0.14	0	-0.08
F170W	14/2	-2.39	-0.71	0	-2.94
F170W	14/2	-2.04	-2.08	0	-3.21

(Note: The positive values in WF2 for the first entries in the table may reflect a slight change in the illumination pattern of the Cal Channel at the 0.1% level.)

amplitude should be less than 1%. Since the worm structure appears to consistently appear in the same locations, it should be possible to construct correction flats to bring them to an even lower level. In the far-UV, the effect is larger, so the field dependent contamination may limit the accuracy of even relative photometry to several percent.

Aside from the “worm” features, there is evidence for variations in contamination as a function of the location in the field. Table 3 shows the mean percentage changes of the ratios of on-orbit internal flats to TV internal flats near the centers of the CCDs, normalized to unity in WF3. These data show that all CCDs and all filters show declining response with time relative to WF3; this is inconsistent with the data shown in Figs. 8(a) and 8(b), which indicate less of a decline in PC1 than in WF3. Some areas in the chips have larger variations than the centers; for example, the pyramid corner of WF4 is down by about 4%–5% in *F170W* and *F255W* relative to WF3.

Since the level of the calibration channel lamps cannot be assumed to be constant, the numbers are only relative to WF3. As such, these numbers represent a lower limit on the total attenuation. If we assume an attenuation of $\sim 10\%$ at *F170W* over the first two months based on the UV star monitor data presented above, these numbers suggest that the ratio of total to relative attenuation is ~ 3 –5.

5.4 QE as a Function of CCD Temperature

Figure 9(b) shows that the apparent throughput increased in most filters when the CCD temperature was changed from -76 to -88 °C. We expected an increase in apparent throughput because of the improvement in the CTE; if the CTE had been removed entirely, we would have predicted an apparent improvement of $\sim 6\%$ in throughput for the flux standards which are measured in the centers of the chips. We see gains of this amplitude in *F439W* and *F555W*, but significantly smaller gains in *F814W*. This suggests that the CCD QE has changed with the CCD temperature and that the QE change is a function of wavelength; the QE appears to be lower in the near-IR at -88 °C than at -76 °C. Analysis of internal flat data appears to support the conclusion that the QE has changed in the near-IR. In the UV, the apparent throughput increased by $\sim 10\%$ [see Fig. 8(a)]; perhaps this change was so large because the CTE effect is larger in the far-UV where the starlight is spread over more pixels (because of scattering in the optics).

5.5 Measurement of Throughput

It is important to determine the observed throughput of the WFPC2 to determine the feasibility of and exposure times for science programs. This has been done using observations of flux standards (AGK+81 266, BD+75 325, GRW+70 5824, Feige 110) through a variety of filters.

Count rates are predicted from sensitivity curves for the telescope, instrument, filters, and detectors. Various predictors are presented in the WFPC2 instrument handbook, the STSDAS package SYNPHOT, the XCAL package, and several software tools used by the WFPC2 IDT. In many cases these give similar predictions, but in others they do not. In particular, predictions for some of the UV filters can differ by tens of percent; this has been tracked down to differences in the extrapolation of the input filter files. Interpolation and sampling errors can be important; note, e.g., that the default wavelength tables used by SYNPHOT and XCAL may not sample the narrow-band filters at enough wavelengths.

Pre-launch estimates of count rates were compared with observations to derive modifications to the input sensitivity curves. Most of the pre-launch estimators did not incorporate throughput information obtained in TV. During TV, we discovered that the throughput in the far UV was roughly 60% of expectation because of lower mirror reflectivities and also found that the filter curve for *F160W* needed to be adjusted. The observed counts rates were compared to our best estimate from pre-launch data, including the modified far-UV response. At visible wavelengths, the pre-launch estimates were accurate to several tens of percent. In the far UV, however, the response was down even further than expected from the TV data, though possibly within the substantial uncertainties. The far-UV response was a factor of 3–5 lower than expected from predictors which did not incorporate the TV data, such as the pre-launch versions of the WFPC2 handbook and SYNPHOT.

The on-orbit observations have been used to derive rough corrections to the sensitivity curves. These corrections have been incorporated into the second version of the WFPC2 handbook and the STScI database. Figures 10(a) and 10(b) present the ratio of the observed-to-predicted count rates using the new curves. The absolute normalization was derived using an estimate of the CTE effect and is only approximate. With the revised curve, the predicted count rates are generally within $\sim 10\%$ – 20% of the observed values; this should be adequate for exposure time prediction.

In principle, it should be possible to use the adjusted throughput curves with filter profiles and input spectra to get a photometric calibration. Figures 10(a) and 10(b) indicate some of the uncertainties with this method. In Fig. 10(a), observations of 3 separate stars are shown, and it is clear that there are systematic differences in results between different stars; a similar effect is seen in the two stars in Fig. 10(b). Absolute spectrophotometry is difficult, so perhaps it is not surprising that there are constant multiplicative shifts between different stars. However, it appears that there are also wavelength dependent differences between stars; note the difference between the circles (AGK+81 266) and the squares (GRW+70 5824) for the *same* set of filters.

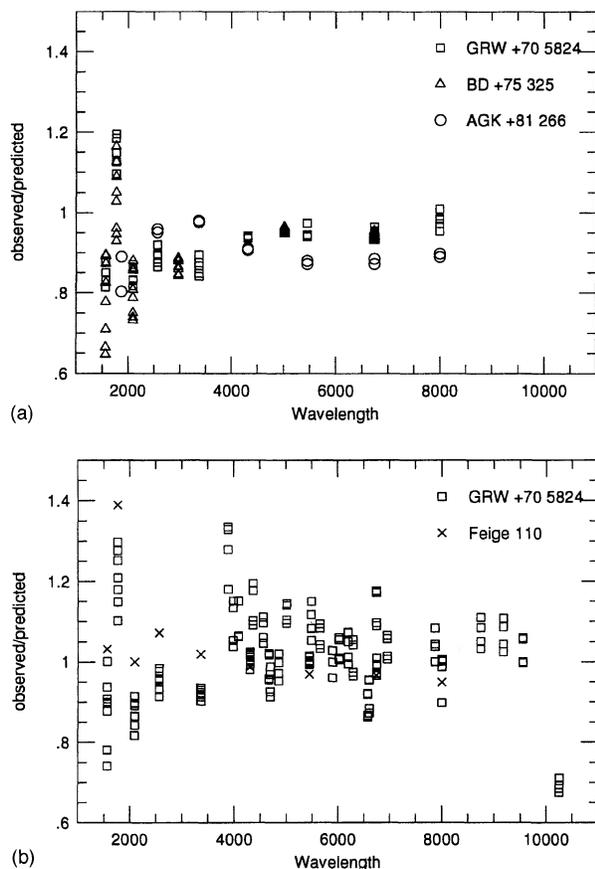


FIG. 10—(a) Comparison of observed throughput with synthetic predictions at -76°C . Different symbols represent different flux standards. (b) Comparison of observed throughput with synthetic predictions at -88°C . Different symbols represent different flux standards.

Figure 10(b) plots data for a large suite of filters. It is apparent that no smoothly varying adjustment will fit the entire set of filters to better than $\sim 5\%$ – 10% . This suggests that there may be errors in the filter database at this level.

Because of these problems, users are cautioned against inferring accurate calibration on the basis of synthetic calculations from throughput data. We are currently attempting to derive the best corrections to the sensitivity curves (e.g., filters, instrument, and CCD) to match observed values. These will be presented in a separate paper on WFPC2 photometric calibration, along with additional discussion of the expected accuracy.

5.6 Calibration and Accuracy of Broad-Band Photometry

Because the filters used on WFPC2 differ significantly from those used for most ground-based work and because of the premium on *HST* observing time, photometric calibration is fairly involved. In this section, we briefly discuss some of the issues involved with WFPC2 photometric calibration, describe the data being taken to address these issues, and estimate the uncertainties we expect to have with various levels of photometric calibration. The results of the photometric calibration will be presented in a separate publication.

In astronomy, much photometric work is done with reference to some established standard photometric system. Physical inferences from data on such a system is drawn either from comparison against a large set of objects in this photometric system (e.g., comparison of a color–magnitude diagram against those of globular clusters of known properties) or by using some physical calibration of this system (e.g., determination of cluster ages and metallicities by taking physical stellar models and estimating how physical quantities map into the standard system). Note that both of these processes require a fair amount of calibration observations, since knowledge of absolute responses of telescopes, filters, and detectors is very difficult to obtain accurately.

We consider three ways in which physical information can be drawn from observations with WFPC2. (1) Transform WFPC2 magnitudes onto an already well-studied photometric system and draw physical inferences on the basis of previous work. (2) Accumulate a database of observations in the WFPC2 system and make physical inferences based on comparison with these observations. Most likely, a large reference set of WFPC2 observations of objects about which much is already known will not take place with *HST* because of the premium on observing time; however, it is possible to approximate the WFPC2 system from the ground and accumulate data there. (3) Attempt to understand the absolute response of WFPC2 and use physical models of the objects being studied.

In practice, all of these approaches will be used, and it is important to understand the limitations of each. Some of the problems with determining the absolute response of WFPC2 were discussed in the previous section; currently these problems limit the accuracy of synthetic photometric calibration to 5%–10%.

Transformations between observed WFPC2 magnitudes and those on a standard system will depend on the spectral shape of the object being observed; the larger the deviations of the WFPC2 response from that of the standard system, the more sensitive the transformations will be to differences in the underlying spectrum. Because of this, transformations to a standard system will only be accurate to the degree to which the spectrum of the object matches the spectrum of the standards. An example of this problem might be that photometric transformations for stars are dependent on the stellar metallicity; if one applies a transformation derived from stars of one composition to stars of a different composition, systematic errors might arise. Similar errors might occur for galaxies and for stars with different reddenings. The magnitude of such errors, based on models and some observations, is expected to be less than 5% in most cases, and will be discussed in more detail in our photometric calibration paper.

To determine transformations, observations must be made of standards over as wide a range of color as those objects being studied. Because there is a limit to how much time *HST* spends on calibration, we have currently only obtained data of two fields with stars with a moderate range of colors (in the outer regions of the globular clusters ω Cen and NGC 6752) and of a few spectrophotometric standards. To allow a sampling of stars of more diverse spectral types, we have observed these fields from the ground with a setup (filters

and detector) which match the flight system as closely as possible. This minimizes errors in the transformations between the flight and ground systems which are determined only over the limited color range available in the flight calibration fields. We have used the ground system to observe stars of a large range of spectral type to determine transformations to a standard system. The ground observations are also being used to define a WFPC2 photometric system. With observations of standard stars, the zeropoints of the new WFPC2 system can be defined to match magnitudes on the standard system for a specified type of star.

This procedure was also performed for the original WF/PC broadband filters (Harris et al. 1993; Harris et al. 1991; WF/PC Science Verification Report) to get transformations to the Johnson-Cousins *UBVRI* system. A WF/PC ground system was defined from observations of standards such that stars with zero color in the *UBVRI* system have equal WF/PC and *UBVRI* magnitudes. However, the WFPC2 photometric system differs from that of WF/PC because of some different filter bandpasses and because the response of the Loral CCDs in WFPC2 differs significantly from the TI CCDs in WF/PC.

Our ground-based data, in combination with observations of flux standards on-orbit, give us information on how to best adjust the instrument/filter throughputs; this should result in improvements in the synthetic photometry, and hence for comparison with physical models. Synthetic results will also probably have to be used for filters in which we have not obtained ground calibration data. The ground-based work will also provide WFPC2 system magnitudes for a set of standard stars, which will be of use to observers trying to match the WFPC2 system from the ground. Results from this work and from the analysis of on-orbit photometric calibration observations will be presented in a separate paper.

5.7 The Effect of Subpixel QE and Undersampling on Photometry

One concern with photometry of WFPC2 images is whether there is a photometric effect from the combination of undersampling of stellar images and subpixel QE variations in the CCDs. To measure the amplitude of such an effect, a series of 20 exposures was taken in ω Cen with dithering of ~ 0.2 WF pixels between each exposure; dithers were made along both rows and columns. At each pointing, several dozen stars in each chip were measured. From the 20 different observations, a mean brightness was found for each star. Figure 11 presents the deviations from these means as a function of pixel centering along both rows (top panel) and columns (lower panel) for stars in WF2 with at least 1000 total DN; the data are binned into 0.1 pixel intervals. There appears to be a systematic error in photometry as a function of the centering along columns at the $\lesssim 2\%$ level; for observations of many objects, this effect would contribute $< 1\%$ rms scatter. Any similar effect along rows is at a significantly smaller amplitude.

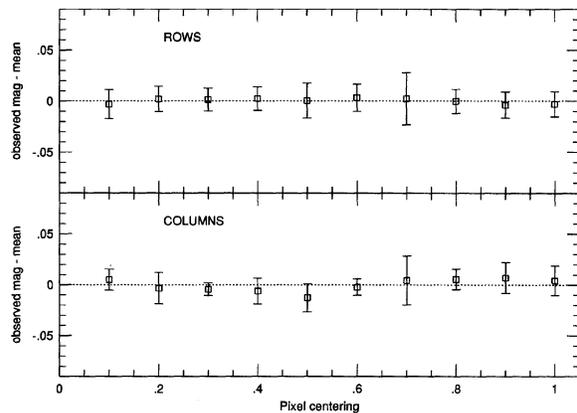


FIG. 11—Average differences between individual magnitudes and their mean for stars in the ω Cen fine dither test as a function of pixel centering.

6. COSMIC RAYS

6.1 Identification and Treatment of Cosmic Ray Events

Two sets of three 2000 s dark exposures were used to determine the properties of on-orbit cosmic ray events for WFPC2. The dark exposures were taken on 1994 January 1 and January 17–19, respectively, with the camera at the -76°C operating temperature. All frames were taken with a gain of $7e^-/\text{DN}$ and all references to DN in this section refer to this gain value. In this analysis, we make no distinction between true cosmic rays and local radioactivity, and we occasionally refer to the total signal collected as the “energy” of an event (which for cosmic rays does not necessarily relate to the energy of the particle). We also have not yet characterized variations in cosmic-ray rates or properties as a function of the *HST* orbit.

The treatment of all dark frames was fully automated to obtain a statistically accurate representation of their properties. All the input frames were first corrected for bias and dark current. Cosmic rays events were identified by looking for pixels with a positive signal more than 5σ (rms) above the mean in at least one pixel. For each identified event, pixels in the neighboring area were examined, and all those with values above a given threshold (2σ above the mean) were assigned to the same cosmic-ray event. The procedure was iterated so that, for each event, all pixels in a contiguous region above the threshold were included. Visual inspection of many events confirms this as a viable procedure to identify the total area affected by each cosmic ray. Because of the large number of events for each frame, occasionally two or more events will be merged into one; we have not corrected for this phenomenon, as it affects a very small number of events. In order to avoid possible edge problems, only the region between pixel 30 and pixel 770 in each coordinate has been considered for each chip, thus including about 86% of the total chip area; corrections for the reduced chip area have been incorporated in the rates given below.

Among the “events” detected are a relatively large (5%–10%) number of single-pixel, low-signal events. The majority of these appear to be just above the detection threshold

(between 5 and 10 DN) and to recur in the same position, and therefore cannot be identified with genuine cosmic ray events, but are probably the product of an inaccurate super-dark or non-Gaussian fluctuations in the dark current rate. For the purpose of this report, we have excluded from the cosmic-ray list a total of 7177 single-pixel events with signal <10 DN. A total of 2543 more single-pixel events remain in the list, some of which may be higher-amplitude fluctuations. Only half of these have a signal larger than 30 DN, while the typical peak signal for multiple-pixel events is around 100 DN, with fewer than 3% of these having a peak below 30 DN. The fraction of these higher-amplitude single-pixel events that recur in the same pixel is also fairly high, but lower than for the low-amplitude events, and probably less than 50%.

While the present data are not conclusive as to the actual rate of occurrence of single-pixel cosmic-ray events, these appear relatively rare in comparison with more energetic ones, at least at or above the 5 DN level. Perhaps less reassuringly, dark current fluctuations seem not uncommon at levels of 5–30 DN. We have not yet performed the analysis for frames taken at -88 °C, where the dark current is significantly lower.

The total number of cosmic rays measured was $\sim 75,000$, corresponding to ~ 3600 CR events/chip/2000 s, or 1.8 events/chip/s. Observers are cautioned to remember to include exposure overheads if this rate is used to estimate the number of cosmic rays in shorter exposures.

6.2 Distribution of Total Signal

For the purpose of signal and size distribution, each event is characterized by its total signal, peak signal, and number of pixels affected.

The histograms in Figs. 12(a) and 12(b) show the distributions of total energy and of energy per pixel of all cosmic ray events, where the total energy is defined as the total number of DN in all pixels affected by an event. One striking feature, already discovered during TV, is the very small number of low-energy events, below about 30 DN. This drop is well above the noise level, and is also clearly above the energy level of excluded single-pixel events.

Just as in TV, the cumulative energy distribution above 100 DN is well described by a Weibull function with exponent 0.25 (Fig. 13). This function has the form

$$N(>E) = N_0 \exp[-\lambda(E^{1/4} - 100^{1/4})],$$

where $N(>E)$ is the number of events with total energy greater than E . A reasonable fit to the distribution of observed events has $N_0 = 2800$ events/chip/2000 s, and $\lambda = 0.92$ (dashed line). Below $E \sim 100$, the cumulative distribution flattens considerably, and the Weibull distribution should not be extrapolated to faint energies. The question of whether substantial numbers of low-signal cosmic rays (≤ 10 DN) exist is not yet settled.

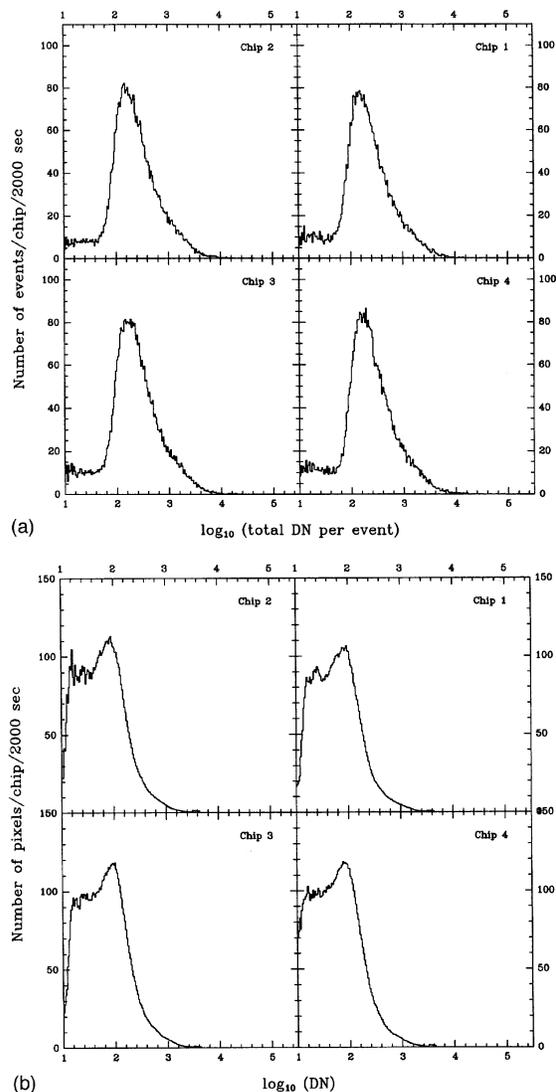


FIG. 12—(a) Histogram of observed cosmic ray event energies. (b) Histogram of number of pixels affected by cosmic rays as a function of DN level.

6.3 Size Distribution

An important characteristic of WFPC2 cosmic-ray events is that they typically affect several pixels. This was already known from TV, and is shown in Fig. 14. The distribution peaks at 4 pixels, and 80% of the events affect between 3 and 9 pixels. Some events have been seen to affect 50 pixels or more. The average number of pixels affected is 6.7 per event, or the equivalent of $\sim 24,000$ pixels per chip per 2000 s exposure. This means that about 3.7% of the pixels will be affected by a cosmic ray in a 2000 s exposure. As cosmic rays can be expected to be randomly placed, about 1000 pixels per chip will be affected in both of two 2000 s exposures; this would make cosmic-ray removal impossible for those pixels. Only about 220 pixels will be affected in both of a pair of 1000 s exposures. Because cosmic rays are multiple pixel events, one would expect that coincident affected pixels will not be randomly placed, but will occur in clumps with a size of a few pixels.

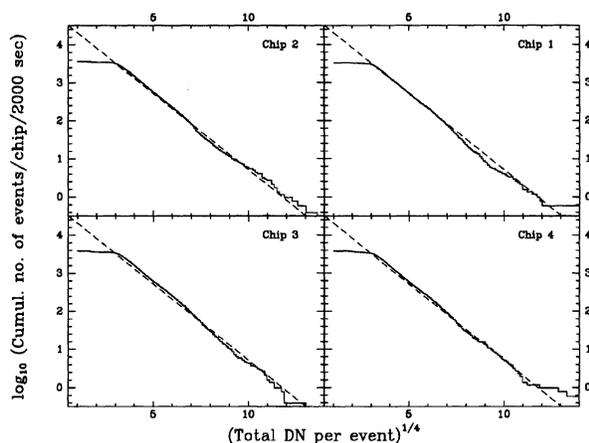


FIG. 13—Cumulative distribution of observed cosmic-ray event energies and a fit of a Weibull function to this distribution.

7. GEOMETRIC DISTORTION IN WFPC2 AND EFFECTIVE PIXEL AREAS

7.1 Geometric Distortion

The WFPC2 cameras are known to have geometric distortion, much of which arises from the presence of the field flattener. This distortion is important for astrometry and for registering images taken at different pointings; it can be a large effect, with “true” positions differing from observed positions by several pixels in the corners of the chips. It is also important to understand the distortion from the point of view of photometry and flat fielding; the distortion causes pixels to have different effective areas as a function of position (see Sec. 3.6.6). This effective area must be accounted for when doing integrated photometry (as opposed to surface photometry).

Since the field flattener is a refractive element, the distortions that it causes are wavelength dependent. The variation with wavelength depends on the index of refraction; for the MgF from which the field flattener is made, the distortion is fairly constant across optical wavelengths (≥ 4000 Å), but

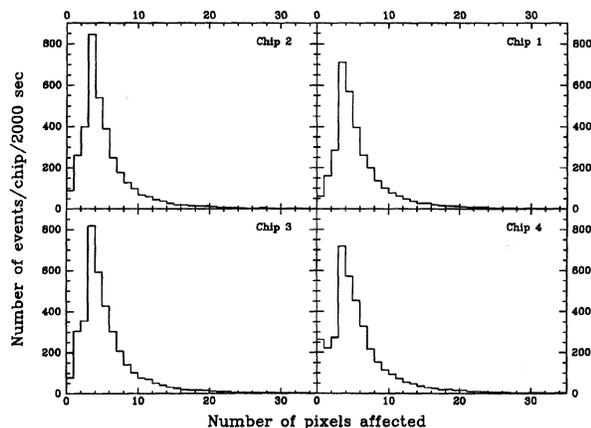


FIG. 14—Histogram of distribution of cosmic-ray event sizes (in pixels).

starts to vary fairly rapidly at shorter wavelengths. This section presents an observational derivation of the distortion around 5500 Å; the wavelength dependence of distortion based on ray-trace calculations using an “as-built” prescription of WFPC2 will be presented in a separate paper. To give a qualitative feeling for the differences in distortion as function of wavelength, we used the results of the ray trace to compute positional differences relative to 5500 Å across a full diagonal of a WF chip as a function of wavelength. We found that this difference is about -0.4 pixels at 9000 Å, about 0.35 pixels at 4000 Å, about 1 pixel at 3000 Å, about 3.5 pixels at 2000 Å, and about 9 pixels at 1500 Å, with positive differences representing larger separations. Clearly, the wavelength dependence cannot be ignored at UV wavelengths.

Note also that because of slightly different wedges in the filters, positions of objects can shift slightly for different filters even at the same wavelength; the shift from filter wedge is just a constant offset across the focal plane.

An estimate of the geometric distortion in the WFPC2 cameras was made from a series of 30 *F555W* exposures in a field in the globular cluster ω Cen. Between each of these exposures, the telescope pointing was shifted by $16''$ in a 6×5 grid along rows and columns. On any given frame, positions of several dozens of stars were measured, amounting to thousands of stellar measurements over the full set of exposures. Distortions were derived from the full set of observations in all four chips simultaneously, solving for distortion, chip rotations, chip offsets, relative positions of all of the stars, and the dither spacing in pixels. For the solution, only stars which appeared on at least 10 frames were considered; this still left ~ 5000 observations of ~ 300 different stars with fairly uniform coverage of the CCDs. Absolute astrometry is not available for this field, so only relative scales, pointings, and rotations were derived. The absolute scales and rotations were derived using additional data (see below).

The dominant distortion is expected to be cubic arising from the field flattener. This distortion should be symmetric around the optical center of each camera. However, not all of the WFPC2 optics are symmetrical around this point because there are some tilts in the system; in addition, the *HST* optical telescope assembly (OTA) can contribute some distortion centered around the OTA axis. Consequently, we use a general cubic distortion representation given by

$$x' = C_{1i} + C_{2i}x + C_{3i}y + C_{4i}x^2 + C_{5i}xy + C_{6i}y^2 + C_{7i}x^3 + C_{8i}x^2y + C_{9i}xy^2 + C_{10i}y^3 \quad (1a)$$

$$y' = D_{1i} + D_{2i}x + D_{3i}y + D_{4i}x^2 + D_{5i}xy + D_{6i}y^2 + D_{7i}x^3 + D_{8i}x^2y + D_{9i}xy^2 + D_{10i}y^3, \quad (1b)$$

where

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

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

This defines the distortion to be zero at pixel location (400,400) apart from chip offsets. We chose the scale and orientation of the center of PC1 to define the global corrected

TABLE 4
 Cubic Distortion Coefficients

Coef	PC1	WF2	WF3	WF4
		before March 4:		
C_1	$3.55437E+02$	$-8.12003E+02$	$-8.06243E+02$	$7.71898E+02$
D_1	$3.44146E+02$	$7.66592E+02$	$-7.70574E+02$	$-7.74071E+02$
		after March 4:		
C_1	$3.54356E+02$	$-8.12003E+02$	$-8.07068E+02$	$7.72904E+02$
D_1	$3.43646E+02$	$7.66592E+02$	$-7.71489E+02$	$-7.74638E+02$
		remaining coefficients are valid all times:		
C_2	$1.00021E+00$	$1.95172E-02$	$-2.18757E+00$	$1.37619E-02$
C_3	$9.79758E-04$	$-2.18805E+00$	$-8.51284E-03$	$2.18899E+00$
C_4	$9.84222E-08$	$-1.51015E-06$	$-3.21458E-07$	$1.33326E-06$
C_5	$-6.31327E-09$	$5.06693E-06$	$2.77459E-06$	$-4.00828E-06$
C_6	$-7.19983E-07$	$1.04063E-06$	$-2.06874E-06$	$-5.71282E-07$
C_7	$-3.73922E-08$	$-4.55094E-10$	$7.39223E-08$	$-2.00520E-09$
C_8	$6.65101E-10$	$7.50249E-08$	$-1.18833E-09$	$-7.86673E-08$
C_9	$-3.51470E-08$	$-1.51652E-09$	$7.73056E-08$	$-3.37491E-09$
C_{10}	$-2.55003E-09$	$7.34568E-08$	$-3.72971E-10$	$-7.81490E-08$
D_2	$1.00544E-03$	$2.18607E+00$	$7.40319E-03$	$-2.18688E+00$
D_3	$9.99786E-01$	$2.02006E-02$	$-2.18554E+00$	$1.29831E-02$
D_4	$-5.80870E-07$	$-3.79469E-06$	$-1.49841E-06$	$1.30814E-06$
D_5	$-5.07221E-07$	$-2.99518E-06$	$2.96650E-06$	$2.01760E-06$
D_6	$3.41574E-07$	$1.05789E-07$	$-2.02756E-07$	$-8.47656E-07$
D_7	$-1.08091E-09$	$-7.47857E-08$	$-1.45720E-09$	$7.70408E-08$
D_8	$-3.42070E-08$	$-1.67666E-09$	$7.72421E-08$	$-1.68514E-09$
D_9	$1.17819E-09$	$-7.55302E-08$	$8.52633E-10$	$7.75257E-08$
D_{10}	$-4.48966E-08$	$3.02579E-10$	$7.68023E-08$	$-3.69210E-10$

coordinate system; we set the origin of the system to approximately correspond to the apex of the pyramid. Consequently, in our global coordinate system PC1 is in the first quadrant, WF2 in the second, etc. Pixels in the global system have the size of PC pixels, and application of the transformation brings positions from all chips into the orientation for PC1. Because we could not solve for absolute scale or orientation, we set the linear coefficients in PC1 to zero for a first pass. However, we discovered that there are small linear terms in the distortions aside from rotation and uniform expansion, so we used the first iteration to give the relative stellar positions, then solved for the distortion again allowing all terms to be fit. The resulting coefficients were then scaled so that Jacobian of the transformation is unity at (400,400) in PC1. The final coefficients are given in Table 4. As a check, we fit a purely radial cubic distortion model to our solution and found a fairly close match. We also obtained similar results to those derived from ray trace analysis.

The rms scatter of the observations around the fit was ≤ 0.25 PC pixels, corresponding to ~ 10 milliarcsec; this is comparable, though slightly larger, than the measurement accuracy in the WF cameras. No systematic trends were seen in the residual maps.

The cubic distortion from the fit is shown schematically in Fig. 15 for PC1; the arrows show the direction of the distortion, not the correction. The distortion is very similar in the other three cameras. The amplitude of the distortion has been increased in the figure for display purposes; in the corners of each chip the distortion has a true amplitude of about four pixels.

Since the solution was derived for all four chips simultaneously, it yielded chip-to-chip pointings for this set of ob-

servations. However, observers should be aware that the relative chip pointings can vary in time. In particular, adjustments to the articulated fold mirrors (AFMs) in PC1, WF3, or WF4 will change the relative pointings. Such an adjustment was made on 1994 March 4, *after* the set of observations from which we have derived the distortion. Fortunately, we can accurately predict the expected movement as a

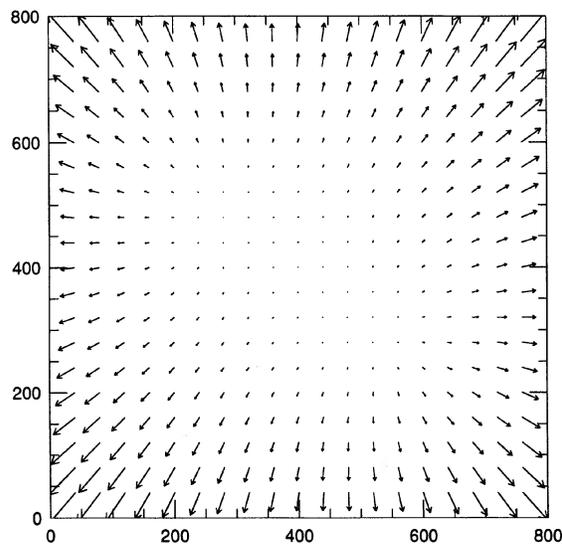


FIG. 15—A map of the distortion in PC1. Arrows show the direction of the distortion; the correction is in the opposite sense. The arrows have been arbitrarily scaled to show the effect; the true amplitude is about four pixels in the corners of the chip. Results are similar for the three other cameras.

result of the AFM motion. The first lines of Table 4 contain the expected offsets for both AFM positions. Observers are cautioned that the positions of the AFMs vary slightly with temperature; consequently, the chip-to-chip pointing may vary by several tenths of a pixel (at maximum, and usually much less) from one observation to another. Because of this, chip-to-chip astrometry will have uncertainties at this level.

From the global fit, we derived relative scales and rotations for the four chips; these data are presented in Table 5. The relative scales were derived for the chip centers. The angles give the rotation of the axes of each chip relative to PC1 in degrees on the plane of the sky with positive rotation from the V3 axis towards the V2 axis. Note, however, that the coefficients presented in Table 4 include the relative scales and rotations; if these coefficients are applied, the resulting coordinates have the scale and orientation of PC1. The numbers presented in Table 5 can be used to bring the undistorted coordinates back to the orientation and scale of each individual chip.

We derived the absolute scale by assuming that the dither was accurately spaced by $16''$. This gave a value of $0.04554''/\text{pixel}$ in the corrected coordinate system defined by Eq. (1) (and consequently, for the center of PC1), which in turn gives scales of 0.09961 , 0.09958 , and $0.09964''/\text{pixel}$ for the centers of WF2, WF3, and WF4. As a check, we used some observations in the cluster M67 for which astrometry was available (Montgomery et al. 1993). We measured 11 stars spread over the four chips, used the distortion model to put them on the global coordinate system, and computed a scale of $0.04555''/\text{pixel}$. This gives us some confidence in both our derived scale as well as the chip-to-chip pointings, since multiple chips were used in the M67 solution. The rms scatter in the distortion fit from the M67 was ~ 40 mas, which is consistent with the fact that the astrometric positions were quoted only to ~ 0.1 arcsec. An additional check on the scale was made from observations in WF4 of several moons of Uranus (Trauger, private communication); the scale derived from these observations agrees extremely well with the value presented above.

We also compared the observed value for the scale with that predicted from ray trace analysis and found rather large disagreement, with the model predicting a scale of $0.09926''/\text{pixel}$ in the WFs; however, in the PC, the model scale was essentially identical to the observed value. We are currently investigating the source of this problem. At this time, we believe that the observationally derived scale above is significantly closer to the truth.

We determined the absolute rotation using the M67 observations. Nominally, the angle between the X axis and the V3 axis in PC1 is 225° (see Fig. 1.1 in the WFPC2 Instrument Handbook). From two epochs of M67 data, we determined that the true position angle of the V3 axis from the X axis is 225.28° . With this knowledge, it is possible to determine where N is in any given image using the value of PA_V3 in the image header, which gives the position angle (from N, measured towards E) of the V3 axis. The formal error in the determination of the angle between the X axis and V3 is $\sim 0.03^\circ$ but there may be some systematic errors in the calculation of PA_V3 (Cox, private communication) which

would affect our derived number. We hope that the measurement of the absolute rotation will be refined in the near future. To bring the undistorted coordinate system into the nominal orientation for PC1, the coordinate system must be rotated by 0.28 degrees counterclockwise in sky orientation or objects must be rotated by 0.28 degrees clockwise. The relative chip rotations are included in the distortion solution but not this absolute rotation.

7.2 Inverse Transformation and WFPC2 Mosaics

One common use of the distortion solution is to create undistorted representations of data and to combine the data from the separate chips into a single geometrically corrected mosaic. To aid in this process, we provide an inverse distortion solution of the form:

$$x_{\text{obs}} = c_{1i} + c_{2i}x' + c_{3i}y' + c_{4i}x'^2 + c_{5i}x'y' + c_{6i}y'^2 + c_{7i}x'^3 + c_{8i}x'^2y' + c_{9i}x'y'^2 + c_{10i}y'^3 \quad (2a)$$

$$y_{\text{obs}} = d_{1i} + d_{2i}x' + d_{3i}y' + d_{4i}x'^2 + d_{5i}x'y' + d_{6i}y'^2 + d_{7i}x'^3 + d_{8i}x'^2y' + d_{9i}x'y'^2 + d_{10i}y'^3 \quad (2b)$$

This transformation allows users to determine the location in the original images from which data should be obtained (e.g., by interpolation) for any given pixel in the corrected image. Since the undistorted coordinate system was defined to have the scale of the PC, the limits in this system for a single WFPC2 pictures are approximately $(-1750:1750, -1750:1750)$. The pixel locations in the raw data in each chip for the corresponding location in the corrected image can be found by using Eq. (2). For most locations, only data from one chip will contribute, but near the chip boundaries, data from multiple chips can contribute to the total light in the corrected image because, as mentioned earlier, the spherical aberration in WFPC2 causes light from some regions of the sky to be imaged onto more than one of the WFPC2 detectors.

If flattened data is used as input to the mosaic, contributions from only one of the chips should be used in the mosaiced image since the flat attempts to correct for the effective vignetting in each chip. However, for cosmetic purposes, we have had more success knitting unflattened data and then flattening it with raw flats mosaiced together in the same way. When knitting unflattened data, contributions from multiple chips can be used in the regions where there is overlap. This provides smooth transitions between chips and seems to work well visually.

The coefficients for the inverse transformation were derived using the distortion solution above to make a set of

TABLE 5
Relative Chip Scales and Rotations

Chip	Relative Scale	Relative Angle
PC1	1.	0°
WF2	2.1872	89.48°
WF3	2.1866	179.79°
WF4	2.1880	270.35°

TABLE 6
Inverse Cubic Distortion Coefficients

Coef	PC1	WF2	WF3	WF4
c_1	4.28529E+01	4.99901E+01	2.91305E+01	4.06423E+01
c_2	1.01839E+00	3.13580E-04	-4.65410E-01	6.89657E-03
c_3	8.19691E-03	4.64781E-01	-2.53882E-03	-4.65172E-01
c_4	-4.07640E-05	-2.44422E-06	-8.01738E-06	-2.85229E-06
c_5	-2.44982E-05	5.28887E-06	-5.00071E-06	5.16479E-06
c_6	-1.48524E-05	-7.39546E-06	-2.92348E-06	-7.91820E-06
c_7	3.85393E-08	7.48860E-11	-3.32528E-09	2.55817E-11
c_8	-8.32054E-10	3.41559E-09	8.12087E-11	-3.50591E-09
c_9	3.65245E-08	-1.64561E-11	-3.50117E-09	-3.28654E-11
c_{10}	2.58167E-09	3.37577E-09	4.45278E-11	-3.47401E-09
d_1	5.33594E+01	2.27199E+01	4.25019E+01	4.60263E+01
d_2	7.27504E-03	-4.64924E-01	-5.86315E-03	4.64946E-01
d_3	1.02076E+00	8.16708E-03	-4.65858E-01	-1.43623E-03
d_4	-1.27390E-05	-7.93511E-06	-2.74006E-06	-8.16812E-06
d_5	-2.37200E-05	4.68366E-06	-5.44820E-06	5.29464E-06
d_6	-4.77746E-05	-2.98490E-06	-8.09822E-06	-2.99135E-06
d_7	1.04890E-09	-3.30055E-09	4.25231E-11	3.51142E-09
d_8	3.55851E-08	-7.84816E-12	-3.49983E-09	-1.12344E-10
d_9	-1.40528E-09	-3.39413E-09	-5.98929E-11	3.55813E-09
d_{10}	4.65028E-08	4.01436E-11	-3.46944E-09	-4.76932E-11

equally sampled points which were then used as input to a solution of the inverse form. As such, the accuracy of the inverse transformation is probably degraded from the original transformation. The coefficients for the inverse transformation are presented in Table 6; note that these coefficients apply for dates later than 1994 March 4.

7.3 Effective Pixel Areas

The cubic distortion coefficients were used to derive effective pixel areas by computing the Jacobian of the coordinate transformation. These pixel area maps are presented in Fig. 16; contours are shown at half percent levels. For mea-

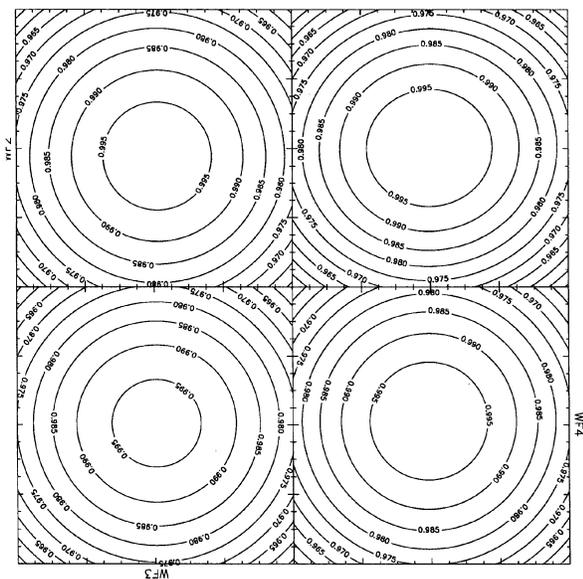


FIG. 16—A map of the effective pixel areas of the WFPC2 chips. The areas are normalized to unity at the center of each chip.

surements of total brightness (as opposed to measurements of surface brightness), data should be multiplicatively corrected by the maps shown in Fig. 16. We suggest that this correction be applied after integrated photometry is measured on flattened frames.

8. CONCLUSIONS

We have presented data describing the performance of WFPC2 on-orbit. Generally, the performance is as expected from pre-launch testing of the instrument. The biggest surprises were the discoveries of a significant CTE problem in the CCDs at an operating temperature of -76°C and a high growth rate of hot pixels at that temperature. However, both of these effects are significantly reduced at the current operating temperature of -88°C . We have derived bias and dark calibration frames from on-orbit data. On-orbit data has validated our flat-fielding strategy; we have derived the illumination pattern from the *HST*+WFPC2 combination and applied this to flat fields obtained during the TV test to produce flat fields which are likely to be accurate to better than a percent on small scales, and better than a few percent across the entire field. The observed PSFs demonstrate that the WFPC2 accurately corrects for the spherical aberration present in the *HST*. We have measured the on-orbit throughput, and find that it appears stable at visible wavelengths; the UV throughput degrades between decontamination procedures, at a rate of about 1 percent per day at 1700 \AA . We are working on accurate photometric calibration, which will be presented in a separate paper.

Our understanding of the WFPC2 and the quality of the calibration products continues to evolve. Consequently, this report only captures the state of our knowledge as of fall 1994. Observers and proposers are cautioned that many of the data presented herein are subject to revision, and are advised to keep abreast of the current knowledge of WFPC2; this information should be available through the STScI.

Calibration data continues to be routinely taken by the STScI. Generally, most of the calibration observations are either internal calibration observations (e.g., biases, darks, and internal flats), external flat-field observations, or routine monitoring of throughput through a subset of filters. There are some additional calibration observations (e.g., photometric measurements through additional filters, observations of an astrometric standard field, etc.) which we hope can be made in the near future; we are working with STScI to implement these. If observers have specific ideas for further calibration work that can be *strongly* justified on the basis of enhanced scientific results, they are encouraged to communicate these to the instrument scientists at STScI.

The WFPC2 IDT is indebted to many people for enabling the instrument to fly and perform superbly. We give special thanks to the entire WFPC2 team at JPL, who built the instrument, to the entire crew of STS61 who accomplished a spectacular repair mission, to the WFPC2 personnel at STScI, and to the WF/PC IDT, who provided large quantities of knowledge and advice throughout the construction and calibration of WFPC2. Among these people, we especially appreciate the comments and work we have received from

Ed Groth, Sylvia Baggett, and Christine Ritchie. This work was supported by NASA under Contract No. NAS7-918 to JPL and subcontract 959145 from JPL to Lowell Observatory.

REFERENCES

- Burrows, C. J., and Krist, J. 1994, Appl. Opt., submitted
Faber, S. M. and Westphal, J. A., eds. 1991, WF/PC Science Verification Report
- Harris, H. C., Baum, W. A., Hunter, D. A., and Kreidl, T. J. 1991, AJ, 101, 677
Harris, H. C., Hunter, D. A., Baum, W. A., and Jones, J. H. 1993, AJ, 105, 1196
Hasan, H., and Burrows, C. J. 1994, PASP, in press
Krist, J., and Burrows, C. J. 1994, WFPC2 Instrument Science Report, STScI
Montgomery, K. A., Marschall, L. A., and Janes, K. A. 1993, AJ, 106, 181
Stetson, P. B. 1987, PASP, 99, 191