The 2005 HST Calibration Workshop

Hubble After the Transition to Two-Gyro Mode

Edited by: Anton M. Koekemoer
Paul Goudfrooij
Linda L. Dressel

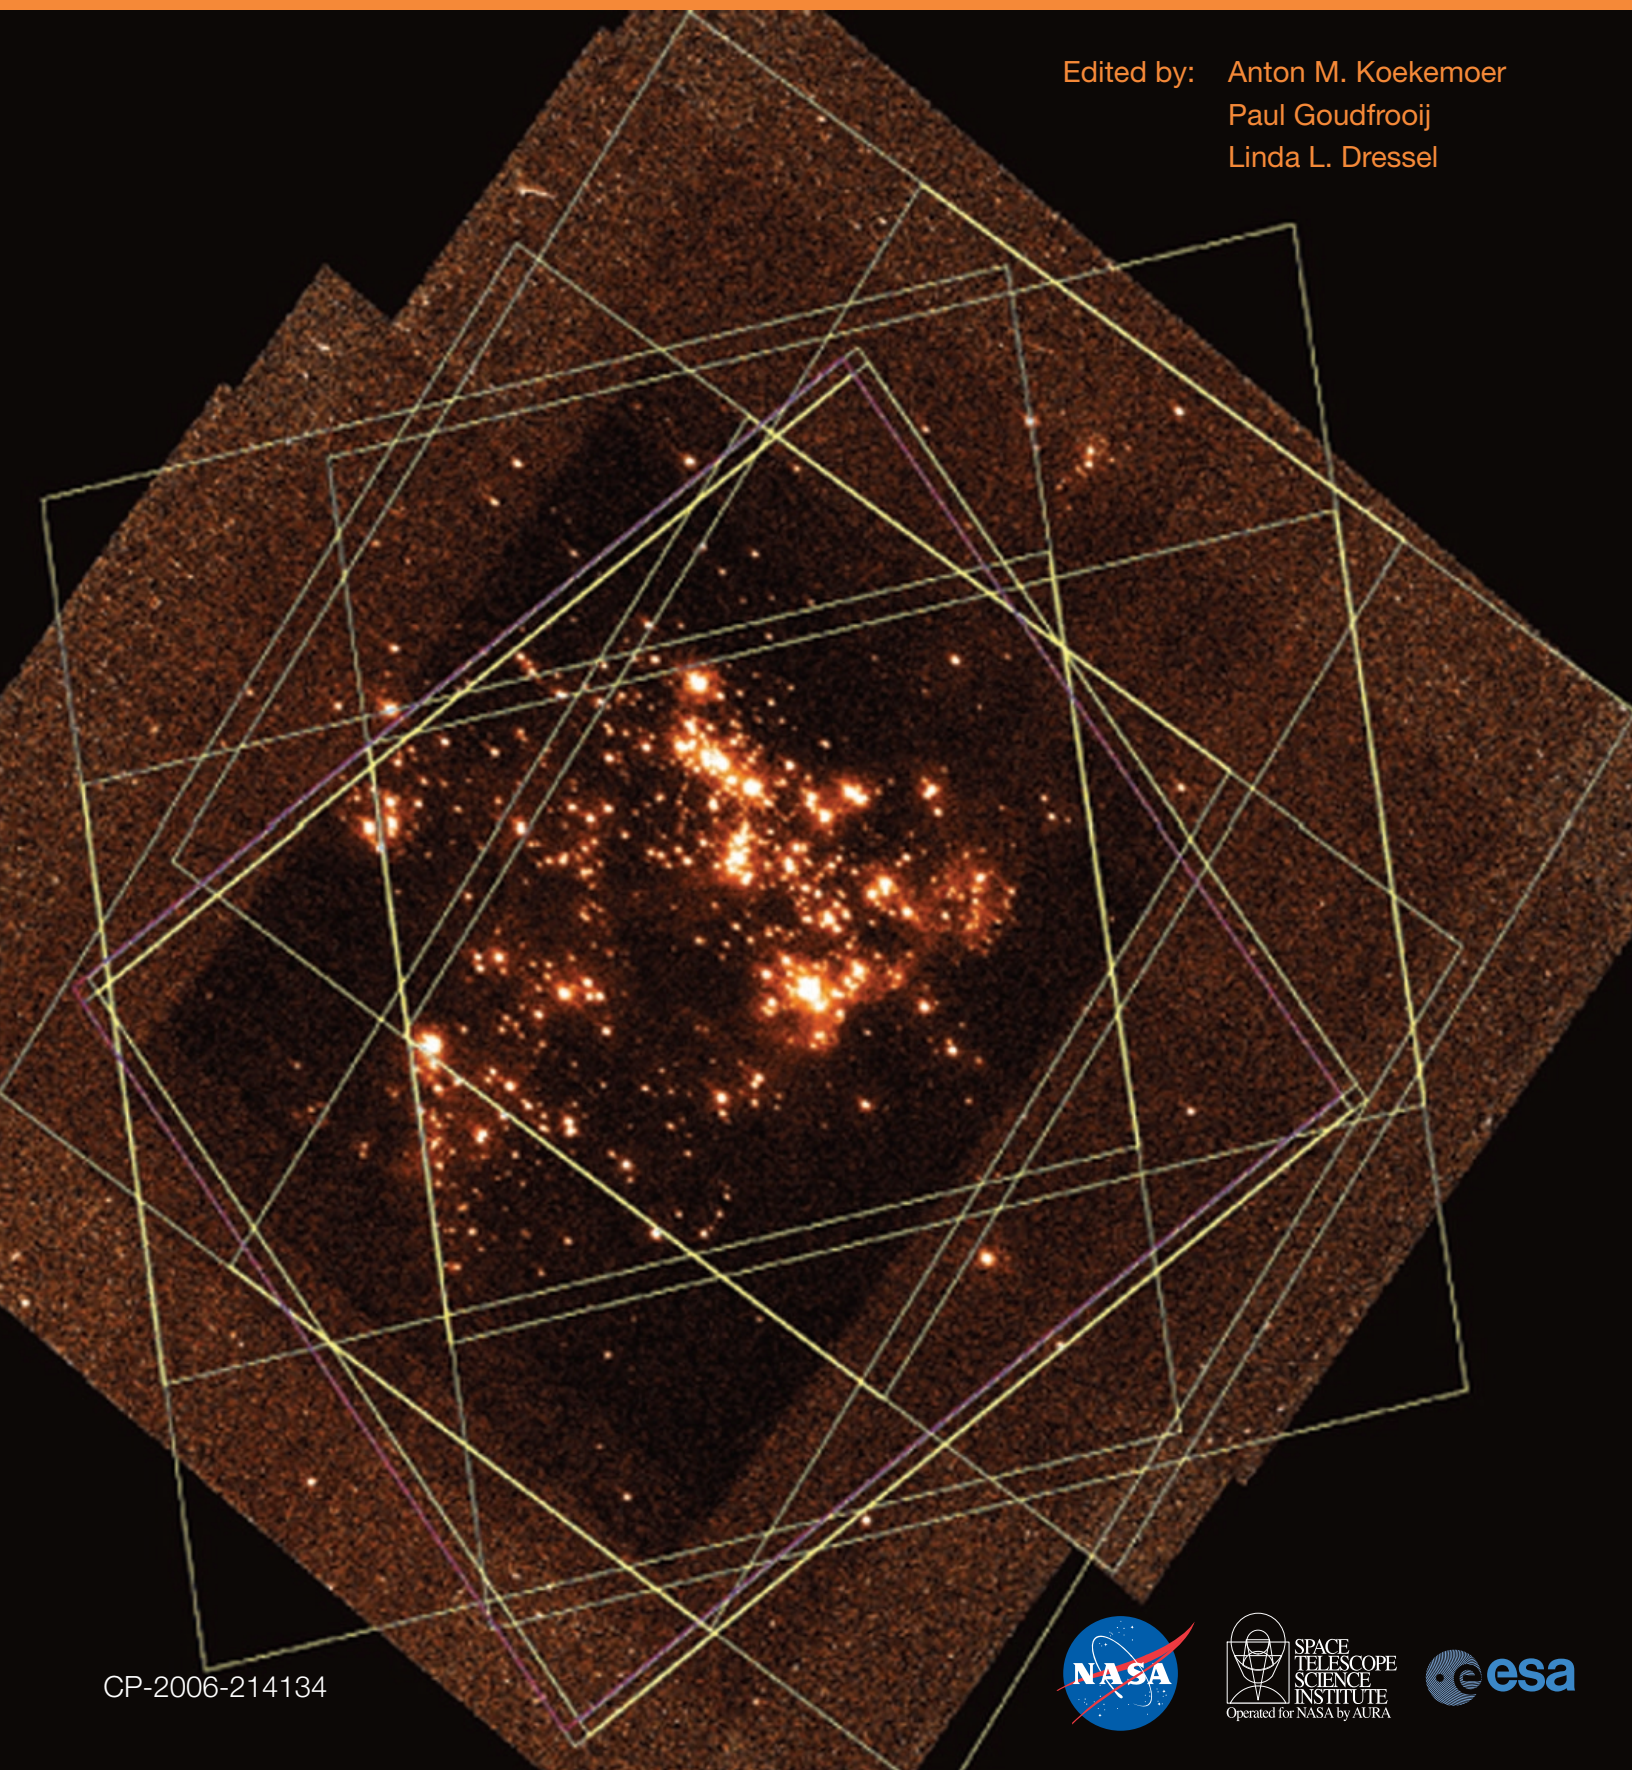

CP-2006-214134

NASA
SPACE TELESCOPE SCIENCE INSTITUTE
Operated for NASA by AURA
esa

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
15-03-2006	Conference Proceedings	Oct. 26–28, 2005

4. TITLE AND SUBTITLE

The 2005 HST Calibration Worskshop
Hubble After the Transition to Two Gyro Mode

5a. CONTRACT NUMBER

NAS5-26555

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Edited by: Anton M. Koekemoer, Paul Goodfrooij, Linda L. Dressel

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Goddard Space Flight Center
Greenbelt, MD 20771

8. PERFORMING ORGANIZATION REPORT NUMBER

2006-00644-0

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSORING/MONITOR'S ACRONYM(S)

11. SPONSORING/MONITORING REPORT NUMBER

NASA CP–2006-214134

12. DISTRIBUTION/AVAILABILITY STATEMENT

Unclassified-Unlimited, Subject Category: 88, 89
Report available from the NASA Center for Aerospace Information,7121 Standard Drive, Hanover, MD 21076. (301)621-0390

13. SUPPLEMENTARY NOTES

A.M. Koekemoer, P. Goodfrooij, and L.L. Dressel: Space Telescope Science Institute/AURA, Baltimore, MD

14. ABSTRACT

The 2005 HST Calibration Workshop was held at the Space Telescope Science Institute during October 26, 2005 to bring together members of the observing community, the instrument development teams, and the STScI instrument support teams to share information and techniques. Presentations included the two-gyro performance of HST and FGS, advances in the calibration of a number of instruments, the results of other instruments after their return from space, and the status of still others which are scheduled for installation during the next servicing mission. Cross-calibration between HST and JWST was discussed, as well as the new Guide Star Catalog and advances in data analysis software. This book contains the published record of the workshop, while all the talks and posters are available electronically on the workshop Web site.

15. SUBJECT TERMS

calibrating, spaceborne telescopes, astronomical catalogs, Hubble

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19b. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			James Jeletic
Unclassified	Unclassified	Unclassified	Unclassified	446	19b. TELEPHONE NUMBER *(Include area code)* (301) 286-6846

The NASA STI Program Office … in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.
- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.
- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.
- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and mission, often concerned with subjects having substantial public interest.
- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at http://www.sti.nasa.gov/STI-homepage.html
- E-mail your question via the Internet to help@sti.nasa.gov
- Fax your question to the NASA Access Help Desk at (301) 621-0134
- Telephone the NASA Access Help Desk at (301) 621-0390
- Write to:
  NASA Access Help Desk
  NASA Center for AeroSpace Information
  7121 Standard Drive
  Hanover, MD 21076–1320

Cover page and artwork designed by Michael Wiggs using a mosaic of an astrometric standard field, NGC 604, made from ACS HRC F220W exposures (Jesús Maíz Apellániz, p. 243). Images of this field and another field (NGC 6681) have been used to derive accurate geometric distortion solutions for the ACS HRC and the STIS NUV-MAMA, and will be used to derive accurate solutions for the ACS SBC (outlines of exposures indicated) and the STIS FUV-MAMA.

NASA/CP2006-214134

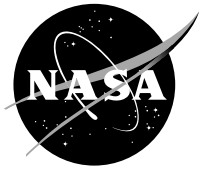

The 2005 HST Calibration Workshop
Hubble after the Transition to Two-Gyro Mode

Proceedings of a Workshop held at the
Space Telescope Science Institute
Baltimore, Maryland
October 26-28, 2005

Edited by Anton M. Koekemoer, Paul Goudfrooij, and Linda L. Dressel

National Aeronautics and
Space Administration

Goddard Space Flight Center
Greenbelt, Maryland 20771

March 2006

Available from:

NASA Center for AeroSpace Information
7121 Standard Drive
Hanover, MD 21076-1320
Price Code: A17

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Price Code: A10

Contents

Part 2. NICMOS

Part 3. STIS, FOS and GHRS

Part 4. Future Instruments: WFC3, COS and JWST

Part 5. Observatory and Two-Gyro Mode

Part 6. GSC-II and Data Analysis

Preface

In 2005, the Hubble Space Telescope (*HST*) celebrated its 15th anniversary of on-orbit operations. Throughout this period, each new generation of instruments has opened up a quantum leap in discovery potential. Most recent is the Advanced Camera for Surveys (ACS), installed in 2002, which has played a prominent role in *HST* science. In addition, the NICMOS Cooling System (NCS) returned the telescope to the forefront of near-infrared astronomy. The next generation of instruments, the Wide Field Camera 3 (WFC3) and the Cosmic Origins Spectrograph (COS), are likewise expected to yield rich discoveries.

Moreover, some steps have been taken to extend the life of *HST*, thereby also extending its scientific potential. Most notable was the transition to Two-Gyro Mode (TGM) in August 2005, activating a new attitude control system that requires only two gyroscopes and allows other gyroscopes to be turned off, to conserve their lifespan for future use.

With new instruments and modes of operation, new calibration challenges also arise, since forefront science demands pushing the telescope and its instruments to their limits. Each new instrument requires extensive calibration before launch and during on-orbit operations. Instruments that have been in operation for a long time permit detailed studies of changes in long-term behavior. We also learn about instrument failure modes and possible ways to repair them, such as the STIS failure. Finally, when instruments are decommissioned, extensive "close-out" calibrations provide a long-lasting scientific legacy.

The 2005 HST Calibration Workshop was held at the Space Telescope Science Institute during October $26-28$, 2005 to bring together members of the observing community, the instrument development teams, and the STScI instrument support teams to share information and techniques. Presentations included the two-gyro performance of *HST* and FGS, advances in the calibration of ACS, NICMOS, STIS, and WFPC2, results on FOS and GHRS after their return from space, and the status of WFC3 and COS which are scheduled for installation during the next servicing mission. Cross-calibration between *HST* and *JWST* was discussed, as well as the new Guide Star Catalog and advances in data analysis software. A total of 105 astronomers attended the workshop which featured 41 talks and 35 posters, as well as splinter groups, demonstrations, and discussions on various topics.

This book contains the published record of the workshop, while all the talks and posters are available electronically on the workshop website.[1] As our knowledge of the instruments continues to improve, the latest information is always available from the STScI website.[2]

We wish to thank all the participants and organizers for making this workshop a success. The workshop was sponsored by the Instruments Division at STScI, and we thank all the members of the Organizing Committee, as well as Dixie Shipley, Susan Rose, Harry Payne, Calvin Tullos, Mike Wiggs, Helmut Jenkner, Margie Cook, and Robin Auer for all their logistical and administrative support in running the workshop and making these proceedings possible. We also thank NASA Headquarters, the *HST* Project at the Goddard Space Flight Center, and the Space Shuttle Program support staff at the Johnson Space Flight Center and Kennedy Space Center for their outstanding support of all *HST* servicing activities.

We would like to dedicate these proceedings to all the astronauts whose heroic commitment to the Hubble mission makes it possible for humankind to continue exploring the universe in this unique way. We all owe a tremendous debt to these brave individuals.

The Editors
Anton M. Koekemoer, Paul Goudfrooij, and Linda L. Dressel
January 2006

[1] http://www.stsci.edu/institute/conference/cal05

[2] http://www.stsci.edu

The Organizing Committee

Santiago Arribas
John Biretta
Linda Dressel
Paul Goudfrooij (Chair)
Tony Keyes
Jessica Kim Quijano
Anton Koekemoer
John MacKenty
Dixie Shipley
Marco Sirianni

Participant List

Javier Alcolea	Observatorio Astronomico Nacional
Alessandra Aloisi	Space Telescope Science Institute/ESA
Jay Anderson	Rice University
Santiago Arribas	Space Telescope Science Institute
Sylvia Baggett	Space Telescope Science Institute
Paul Barrett	Johns Hopkins University
Dan Batcheldor	Rochester Institute of Technology
Stephane Beland	University of Colorado
Louis Bergeron	Space Telescope Science Institute
John Biretta	Space Telescope Science Institute
Howard Bond	Space Telescope Science Institute
Art Bradley	Spacecraft System Engineering Services
Paul Bristow	ST-ECF
Loren Bruns	Reed College
Howard Bushouse	Space Telescope Science Institute
Stefano Casertano	Space Telescope Science Institute
Marco Chiaberge	Space Telescope Science Institute
Colin Cox	Space Telescope Science Institute
Kris Davidson	University of Minnesota
Kyle Dawson	Lawrence Berkeley National Lab
Roelof de Jong	Space Telescope Science Institute
Guido De Marchi	European Space Agency
Rosa Diaz-Miller	Space Telescope Science Institute
Andrew Dolphin	Steward Observatory
Rodger Doxsey	Space Telescope Science Institute
Linda Dressel	Space Telescope Science Institute
Vitaliy Fadeyev	LBNL
Jacqueline Faherty	SUNY Stony Brook
David Floyd	Space Telescope Science Institute
Scott Friedman	Space Telescope Science Institute
Andrew Fruchter	Space Telescope Science Institute
Ronald Gilliland	Space Telescope Science Institute
Paul Goudfrooij	Space Telescope Science Institute
James Green	University of Colorado
Perry Greenfield	Space Telescope Science Institute
Michael Gregg	U. California, Davis and Lawrence Livermore National Lab
Theodore Gull	NASA/GSFC/EUD
Robert Hanisch	Space Telescope Science Institute
Reinhard Hanuschik	European Southern Observatory
Robert Hill	SSAI/GSFC
Dean Hines	Space Science Institute
John Irwin	SLAC
Kazunori Ishibashi	MIT Kavli Institute
Inese Ivans	Carnegie Observatories
Andreas Kaufer	European Southern Observatory
Florian Kerber	ST-ECF
Charles (Tony) Keyes	Space Telescope Science Institute
Jessica Kim Quijano	Space Telescope Science Institute
Randy Kimble	NASA Goddard Space Flight Center

Anton Koekemoer	Space Telescope Science Institute
Gerard Kriss	Space Telescope Science Institute
Martin Kuemmel	Space Telescope - European Coordinating Facility
Matt Lallo	Space Telescope Science Institute
Wayne Landsman	SSAI
Soeren Larsen	ESO / ST-ECF
Tod Lauer	NOAO
Karen Levay	Space Telescope Science Institute
Ray Lucas	Space Telescope Science Institute
Jesús Maíz Apellániz	Space Telescope Science Institute
Jennifer Mack	Space Telescope Science Institute
John MacKenty	Space Telescope Science Institute
Russell Makidon	Space Telescope Science Institute
Sangeeta Malhotra	Space Telescope Science Institute
Phil Marshall	KIPAC
Andre' Martel	Johns Hopkins University
John Martin	University of Minnesota
Brian McLean	Space Telescope Science Institute
Matthew McMaster	Space Telescope Science Institute
Gerhardt Meurer	The Johns Hopkins University
Andrew Michael	RIT
Bahram Mobasher	Space Telescope Science Institute
Max Mutchler	Space Telescope Science Institute
Gillian Nave	National Institute of Standards and Technology
Edmund Nelan	Space Telescope Science Institute
Jeff Nesbitt	Space Telescope Science Institute
C. R. O'Dell	Vanderbilt University
Steve Osterman	University of Colorado
Nino Panagia	ESA/STScI
Cheryl Pavlovsky	Space Telescope Science Institute
Chien Peng	Space Telescope Science Institute
Steven Penton	University of Colorado - CASA
Eric Perlman	University of Maryland, Baltimore County
Ruth Peterson	UCO/Lick and Astrophysical Advances
Charles Proffitt	Space Telescope Science Institute/CSC
Thomas Puzia	Space Telescope Science Institute
Sun Hong Rhie	None
James Rhoads	Space Telescope Science Institute
Jason Rhodes	JPL
Harvey Richer	University of British Columbia
Barry Rothberg	Space Telescope Science Institute
Elena Sabbi	Space Telescope Science Institute
Kenneth Sembach	Space Telescope Science Institute
Marina Shmakova	SLAC/Stanford
Marco Sirianni	ESA/STScI
Tammy Smecker-Hane	University of California, Irvine
Ed Smith	Space Telescope Science Institute
David Soderblom	Space Telescope Science Institute
Randal Telfer	Orbital Sciences / GSFC
Rodger Thompson	Steward Observatory, University of Arizona
Jeremy Walsh	ST-ECF, European Southern Observatory
Richard L. White	Space Telescope Science Institute
Jim Whiting	Space Telescope Science Institute

Tommy Wiklind	ESA/STScI
Chun Xu	Space Telescope Science Institute
David Zurek	American Museum of Natural History

Part 1. ACS and WFPC2

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Calibration Status of the Advanced Camera for Surveys

K. Sembach

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. The Advanced Camera for Surveys is producing spectacular science results that require careful instrument calibration. This article provides an overview of recent calibration highlights and calibration plans for Cycle 14.

1. Introduction

The Advanced Camera for Surveys (ACS) is operating nominally at the time this review is being written. ACS instrument performance with HST in two-gyro mode is indistinguishable from its performance in three-gyro mode (see Sembach et al., this volume and references therein). The ACS Calibration Team at STScI and ST-ECF continually monitors instrument performance through its ongoing instrument calibration activities, and updates the *ACS Instrument Handbook* (Gonzaga et al. 2005) released with the *HST Call for Proposals* at the beginning of each observing cycle.

Cycle 14 brings some new challenges for the calibration of the instrument, as there is somewhat increased demand for Solar-Blind Channel (SBC) observations and prism/grism observations now that STIS is no longer operating. On behalf of all the people involved in ACS calibration activities, this article contains a brief overview of the instrument status and calibration program for Cycle 14. More information about specific calibrations can be found in accompanying articles by ACS Calibration Team members and Guest Observers in the astronomical community. Instrument Science Reports (ISRs) that provide detailed descriptions of the calibration procedures and analyses can be found on the ACS ISR web page:

http://www.stsci.edu/hst/acs/documents/isrs

Information about calibration activities in previous cycles, including the titles and program numbers of calibration programs, can be found at

http://www.stsci.edu/hst/acs/analysis/calib_plan

In preparation for Cycle 14, the ACS Calibration Team reviewed previous Cycle 13 calibration programs and identified activities that needed to be continued into Cycle 14. It also identified new calibration needs that were not covered by the existing programs. As part of this exercise, the team trimmed the sizes of some routine calibration programs by consolidating and reducing the cadence of observations whenever possible. Table 1 contains a summary of the number of ACS calibration orbits allotted for Cycle 14. With contingency, the total orbit request for the primary calibration program is 1320 orbits (67 external, 1253 internal).

For comparison, the Cycle 13 ACS calibration program required 70 external orbits and 1138 internal orbits. Thus, the number of orbits required for the Cycle 14 calibration program is comparable to that in Cycle 13. Routine monitoring programs take less time than before, but special calibrations added in this cycle require more orbits than the special programs in the previous cycle (see Tables 2-4 in the following sections). A small amount of

Table 1: Cycle 14 ACS Calibration Orbit Summary

Cycle 14	Orbits		
	External	Internal	Outsourced[a]
Total Time (routine + special programs)	62	1142	11
With Contingency	67	1253	11

[a]Two programs (PI = Dolphin): ACS Zero-Point Verification (#10621) and ACS Photometric Calibration from Faint Standards (#10622).

contingency time (8-10%) is included in the calibration program for each cycle to account for unforeseen calibration needs.

The 11 outsourced orbits awarded in Cycle 14 by the HST Time Allocation Committee are included in Table 1 for completeness. Calibration outsourcing is an important component of the calibration program since it allows Guest Observers to obtain calibrations that are not otherwise covered by the primary calibration plan defined by the STScI ACS Team. We encourage Guest Observers to propose calibration GO programs when they need more precise calibration data than the standard instrument calibration activities provide, or when they need calibrations for an uncalibrated mode of operation. Requests and suggestions for calibrations to include in the standard ACS calibration plan are also welcomed and will be given full consideration within the context of the overall calibration plan for the subsequent cycle of observations.

2. Cycle 14 Routine Monitoring Programs

Various routine monitoring and maintenance programs are required to track the performance of the instrument and ensure that observers have the most up-to-date calibrations available to reduce their data. Examples of routine calibrations for Cycle 14 include maintaining accurate flat field, dark, and bias frames, monitoring photometric throughput and stability, and tracking CCD charge transfer efficiency (CTE). A summary of the routine calibration programs for Cycle 14 is given in Table 2, which lists the proposal number, principal investigator, program title, frequency of observations, and the number of external and internal orbits required during the cycle. These programs are described briefly below.

Table 2: Cycle 14 Routine Monitoring Programs

Program ID	PI	Title	Frequency	Orbits		Note
				Ext.	Int.	
10729	Sirianni	CCD Daily Monitor	4/week	0	840	Dark, bias creation
10730	Chiaberge	External CTE Monitor	6 months	9	0	CTE loss calibration
10732	Mutchler	Internal CTE Monitor	yearly	0	35	Check against ground
10733	Cox	CCD Hot Pixel Annealing	monthly	0	143	Includes monthly CTE
10734	Cox	CCD Post-Flash Verification	yearly	0	4	Tracks capability only
10735	Cox	SBC MAMA Recovery	as needed	0	4	After irregular safing
10736	Maiz	UV Contamination Monitor	6 months	4	2	SBC, HRC tracking
10737	Mack	CCD Stability Monitor	3 months	13	0	L-flat, distortion, photometry
10738	Mack	Earth Flats	weekly	0	52	Tracks coronagraphic spot
10739	Bohlin	Internal Flat Fields	4 months	0	44	SBC components once
10740	Bohlin	Photo-Spectrophot. Abs. Cal.	yearly	7	0	Filter throughputs, QE
Total (Cycle 14 routine monitoring programs)				33	1124	

2.1. Darks and Biases

The CCD Daily Monitor program (#10729) measures the read-noise and dark current in the ACS CCDs as a function of time. It also tracks the growth of hot pixels. `CALACS` uses the reference files generated by this calibration program for every CCD science exposure it processes, so it is essential that the properties of the darks and biases be updated frequently. As a result, the number of orbits required for this program accounts for the vast majority of ACS internal calibration orbits. For Cycle 14, the default gain setting for the Wide-Field Channel (WFC) has been changed to `GAIN=2` e^-/ADU, so we will be obtaining dark and bias data to support this new gain setting as well as other supported gain settings. See Lucas et al. (this volume) and Mutchler et al. (this volume) for more information on the automated procedures used to process ACS dark and bias frames for use in `CALACS`. Recent relevant ISRs include:

- *SBC Dark and Cumulative Images*, ACS ISR 2004-14, by C. Cox
- *Bias and Dark Calibration of ACS Data*, ACS ISR 2004-07, by M. Mutchler et al.

2.2. CCD Hot Pixel Annealing

The CCD Hot Pixel Annealing program (#10733) reduces the number of non-permanent hot pixels on the ACS CCDs by warming the detectors to temperatures of +20° C for 6 hours before returning to the normal operating temperatures (-77° C for the WFC and -80° C for the HRC). These anneal times are a factor of two shorter than in previous cycles; program #10453 in Cycle 13 demonstrated the effectiveness of reducing the anneal time from 12 hours to 6 hours. Reduced anneal times make scheduling easier. At the current hot pixel threshold of 0.08 e^- pix^{-1} sec^{-1}, the anneals typically eliminate $\sim$ 82% of the non-permanent hot pixels on the WFC and $\sim$ 87% on the HRC (see Sirianni et al., this volume, for more information). Data from this program are also used in the monitoring of CTE (internal) and the tracking of dark levels.

2.3. Flat Fields

Several routine monitoring programs provide information about the flat field properties of the ACS CCDs as measured using both external observations (stellar photometry and bright Earth illumination) and internal lamp exposures. The CCD Stability Monitor program (#10737) is the primary source of low-frequency flat fields (L-flats). Regular observations of the same star field within the globular cluster 47 Tuc provide L-flats that have an accuracy of $\sim$ 1% over the full fields of view of the HRC and WFC. The observations track changes in relative sensitivity with an accuracy of $\sim$ 0.1% per year. This program also monitors variations in the CCD geometric distortion corrections with time.

The Earth Flats program (#10738) cross-checks the L-flats created by the CCD Stability Monitor and tracks any changes on time scales shorter than the three month interval of the CCD Stability Monitor by obtaining exposures of the bright Earth during occultations. This program also determines the coronagraphic spot position to better than 1 pixel accuracy.

The Internal Flat Field (#10739) program uses the internal tungsten lamp to assess the stability of the pixel-to-pixel flat fields (P-flats) in several HRC/WFC filters (F435W, F625W, F814W). In Cycle 14, new P-flats will be obtained for the six SBC filters and the SBC prisms (PR110L and PR130L). In all cases, the goal is to produce P-flats accurate to better than 1%.

We comment on flats for polarization and ramp filter observations below in §§5 and 6. Further information about flat fields can be found in the article by Mack et al. (this volume). Recent ISRs related to routine flat field creation and monitoring include:

- *SBC L-Flat Corrections and Time-Dependent Sensitivity*, ACS ISR 2005-13, by J. Mack et al.
- *Earth Flats*, ACS ISR 2005-12, by R. Bohlin et al.
- *The Internal CCD Flat Fields*, ACS ISR 2005-09, by R. Bohlin & J. Mack
- *Flats: SBC Internal Lamp P-Flat*, ACS ISR 2005-04, by R. Bohlin & J. Mack
- *ACS Coronagraphic Flat Fields*, ACS ISR 2004-16, by J. Krist

2.4. Charge Transfer Efficiency

The ACS CCD detectors degrade with time due to radiation damage. One effect of this degradation is a decrease in CTE. Several monitoring programs track the level of CTE and its change with time. The External CTE Monitor (#10730) obtains images of 47 Tuc with half-field of view dithers to estimate photometric losses due to CTE as a function of time and position on the CCDs. The goal of this program is to provide corrections for photometric measurements to an accuracy of 1-2%. An independent check on the CTE is provided by data obtained as part of the CCD Stability Monitor program. The ACS Team is currently investigating procedures for mitigating the effects of CTE in the future, including changing the temperature at which the CCDs operate (Cycle 14 special program #10771).

The Internal CTE Monitor (#10732) uses internal tungsten lamp data to trend the overall CCD radiation damage. It tracks both parallel and serial CTE performance. The data for this program are not used directly in the calibration of ACS observations.

The CCD Post-Flash Verification program (#10734) occasionally tests the ability to illuminate the ACS CCDs with a light emitting diode in a repeatable fashion. This capability may be needed in the future to improve CTE for some exposures (at the expense of adding noise to the data). The post-flash is not yet necessary or available for science observations.

A description of the results from the ACS CTE monitoring programs is given by Chiaberge et al. (this volume). Recent ISRs related to CTE include the following:

- *Internal Monitoring of ACS Charge Transfer Efficiency*, ACS ISR 2005-03, by M. Mutchler
- *Time Dependence of ACS WFC CTE Corrections for Photometry and Future Predictions*, ACS ISR 2004-06, by A. Riess & J. Mack
- *Elevated Temperature Measurements of ACS Charge Transfer Efficiency*, ACS ISR 2004-04, by M. Mutchler & A. Riess

2.5. Photometric/Spectrophotometric Throughput and Contamination Monitoring

The absolute sensitivity and repeatability of ACS photometric and spectrophotometric observations are calibrations that require monitoring on a yearly timescale. The Photo-Spectrophotometric Absolute Calibration monitor (program #10740) establishes the relative magnitudes of three primary white dwarf calibrators to 0.1% accuracy and checks repeatability of the WFC and HRC filter throughputs to 0.2% accuracy using observations of single-star flux standards. These measurements are needed to refine the filter bandpasses. A portion of the time in this program is also being used to cross-calibrate the ACS prism, grism, and F850LP filter with NICMOS and STIS. Changes in sensitivity and the CCD quantum efficiency as a function of time are are also tracked by the CCD Stability Monitor (program #10737), which observes large numbers of stars in 47 Tuc.

A comprehensive paper describing the ACS photometric calibration and the photometric transformations to other photometric systems has been completed recently by Sirianni et

al. (2005). The paper includes transformation coefficients for converting ACS HRC/WFC photometry to WFPC2 and the Landolt UBVRI photometric systems. It also contains information about ACS aperture corrections for point source photometry.

The UV Contamination Monitor program (#10736) tracks the throughputs in the six SBC filters (F115LP, F122M, F125LP, F140LP, F150LP, F165LP), the SBC prisms (PR110L, PR130L), three HRC UV filters (F220W, F250W, F330W), and HRC PR200L. These results are cross-referenced to previous STIS observations of the same cluster (NGC 6681). The goal is to track the UV photometry to 1% accuracy to monitor small changes in the throughput on timescales of ~ 6 months.

Recent ISRs related to ACS sensitivity and UV contamination monitoring include:

- *SBC L-Flat Corrections and Time-Dependent Sensitivity*, ACS ISR 2005-13, by J. Mack et al.
- *The Photometric Stability of ACS: Revisiting the Hubble Deep Field*, ACS ISR 2004-17, by A. Riess
- *Detector Quantum Efficiency and Photometric Zero Points of the ACS*, ACS ISR 2004-08, by G. De Marchi et al.
- *Results of UV Contamination Monitoring of the ACS*, ACS ISR 2004-05, by F. Boffi et al.

2.6. SBC MAMA Recovery

The SBC MAMA Recovery program (#10735) is used to turn on the ACS MAMA and return it to its normal operational state after an anomalous shutdown. This program is invoked only when needed (less than once per year).

3. Cycle 14 Special Calibration Programs

Several special calibration programs complement the Cycle 14 routine calibration programs described above. These special programs, which are listed in Table 3, provide basic calibrations for the SBC and ramp filters, information about the UV narrow-band red leak, and the dependence of CTE and QE on CCD temperature. Results from this latter test (program #10771) will be used in assessing the need for installation of the Aft Shroud Cooling System during the next Hubble servicing mission.

Table 3: Cycle 14 Special Calibration Programs

Program ID	PI	Title	Orbits		Note
			Ext.	Int.	
10722	Maiz	SBC Geometric Distortion	6	4	Basic calibration
10731	Chiaberge	UV Narrow-Band Red Leak	2	0	Responds to early failed cal
10741	Suchkov	Continuum L-Flats (Ramps)	3	0	Basic calibration
10742	Fruchter	Ramp, Grism Wavelengths	4	0	Responds to early failed cal
10743	Larsen	Improved Wavelengths (SBC Prism)	2	2	QSO Lyα lines (1400-1800Å)
10771	Sirianni	CTE/QE Temperature Dependence	12	12	ASCS support test
Total (Cycle 14 special programs)			29	18	

A few special programs from Cycle 13 were completed recently; they are listed in Table 4. Some of these were functional checks for capabilities to be used only if needed (programs #10449, 10450), and others improved calibrations for the polarizers and prisms

(programs #10378, 10391). The Short Annealing Test (program #10453) verified the effectiveness of using 6 hour anneals (see §2.2).

Table 4: Cycle 13 Special Calibration Programs

Program ID	PI	Title	Orbits		Note
			Ext.	Int.	
10378	Biretta	Polarization Calibration	12	0	Last visit in July 2005
10391	Larsen	Wave, Flux for Prisms (SBC and HRC)	1	11	Last visit in August 2005
10449	Cox	SBC Filter Wheel Checkout	0	7	Use program only if needed
10450	Sirianni	Functional Test of MEB2 Switch	0	2	Use program only if needed
10453	Sirianni	Short Annealing Test	0	0	Adopted for routine use
10720	Riess	Monochromatic PSF in the Red	1	0	Recently submitted
Total (Cycle 13 special programs)			13	20	

4. SBC Calibrations

SBC observations currently account for only $\sim 3\%$ of the ACS observing time, but observations in this channel have received more attention since STIS shut down. The Cycle 14 SBC Geometric Distortion program (#10722) is intended to improve the geometric distortion solution in the ACS/SBC imaging modes and the PR130L prism mode. The geometric distortion correction is needed to upgrade the SBC L-flats to a level approaching that of the CCD L-flats. Several special calibration programs designed to improve the SBC prism wavelengths (#10743, 10391) enhance previous calibrations. In Cycle 14, prism observations account for about one quarter of the SBC observing time. See the accompanying articles in this volume by Cox and Larsen for more information about SBC calibration activities.

5. Polarized Filters

Polarization calibration activities (such as Cycle 13 program #10378) have been used to characterize the ACS polarizers and their possible uses. A series of ISRs describing the polarizers describes recent progress in calibrating the polarized filters.

- *ACS/HRC Polarimetry Calibration IV. Low-Frequency Flat-Fields for Polarized Filters*, ACS ISR 2005-10, by V. Kozhurina-Platais & J. Biretta
- *ACS/HRC Polarimetry Calibration III. Astrometry of Polarized Filters*, ACS ISR 2004-11, by V. Kozhurina-Platais & J. Biretta
- *ACS Polarization Calibration II. The POLV Filter Angles*, ACS ISR 2004-10, by J. Biretta & V. Kozhurina-Platais
- *ACS Polarization Calibration I. Introduction and Status Report*, ACS ISR 2004-09, by J. Biretta et al.

6. Grism/Prism Spectroscopy and Ramp Filters

In Cycle 14, prism/grism spectroscopy accounts for approximately one tenth of the WFC observing time and one-quarter of the HRC and SBC observing time. The ST-ECF is

responsible for the calibration of the ACS spectroscopic modes. Work continues on characterizing both the wavelength solutions and sensitivities of these modes. Recent ISRs include:

- *Updated Wavelength Calibration for the WFC/G800L Grism*, ACS ISR 2005-08, by S. Larsen & J. Walsh
- *Flat-Field and Sensitivity Calibration for ACS G800L Slitless Spectroscopy Modes*, ACS ISR 2005-02, by J. Walsh & N. Pirzkal

The ST-ECF spectral extraction software, `aXe`, is the primary software used to manipulate ACS slitless spectroscopy images. The most current version of this software (v1.5) can be linked to from the ACS web page. See the accompanying article in this volume by Walsh et al. for a description of the software and its uses.

Two Cycle 14 special programs will provide updated calibrations for the ACS ramp filters. Program #10742 will calibrate the throughputs of the ramp filters as a function of wavelength by obtaining observations of a flux standard with the filters crossed with the grism. This should also calibrate the zeroth order of the grism, which may allow some users to avoid having to obtain additional direct images of the fields they observe with the grism. Program #10741 will provide continuum L-flats for the ramp filters, which will improve the characterization of the total filter transmissions.

7. Coronagraphy

Coronagraphic observations account for roughly 3-4% of ACS observing time. They depend strongly on the pointing stability of the telescope and the ability to flat field detector artifacts. These issues are discussed in the following two ISRs. Observers should note that there is no appreciable difference in the quality of coronagraphic observations between two-gyro mode and three-gyro mode.

- *ACS Coronagraph Performance in Two-Gyro Mode*, ACS ISR 2005-05, by C. Cox & J. Biretta
- *ACS Coronagraphic Flat Fields*, ACS ISR 2004-16, by J. Krist

8. Point Spread Function

The ACS point spread function (PSF) depends on wavelength. To date, there has been only a limited amount of information about the PSF shape. Encircled energies are given by Sirianni et al. (2005), and the shape and stability of the PSF has been documented in the following ISRs:

- *Characterization of the ACS/HRC Point Spread Function in Two-Gyro Mode*, ACS ISR 2005-11, by M. Sirianni et al.
- *Two-Gyro Pointing Stability of HST Measured with ACS*, ACS ISR 2005-07, by A. Koekemoer et al.
- *Multi-Filter PSFs and Distortion Corrections for the HRC*, ACS ISR 2004-15, by J. Anderson & I. King

Several talks at this workshop also addressed the issue of ACS PSFs.

The ACS Calibration Team is using Cycle 14 special calibration program #10720 to characterize the monochromatic PSF of the WFC in the red using a combination of the F850LP filter and several ramp filters. The purpose of this program is to improve the precision of photometric measurements obtained with the F850LP filter.

9. Astrometry

With the advent of Guide Star Catalog 2, it should be possible to improve the absolute astrometric solutions for ACS images by roughly an order of magnitude over current solutions. An accuracy in the range of 0.1-0.3″may be achievable for fields containing sufficient numbers of identifiable guide stars. The following ISR documents the technique for achieving an astrometric accuracy comparable to the resolution of the telescope.

- *Demonstration of a Significant Improvement in the Astrometric Accuracy of HST Data*, ACS ISR 2005-06, by A. Koekemoer et al.

10. WFC CCD Gain Change for Cycle 14

The new default gain setting for the WFC in Cycle 14 is `GAIN=2`. This change helps to alleviate image ghosts caused by electrical cross-talk, which is documented in the following ISRs:

- *Cross-Talk in the ACS WFC Detectors II. Using* `GAIN=2` *to Minimize the Effect*, ACS ISR 2004-13, by M. Giavalisco
- *Cross-Talk in the ACS WFC Detectors I. Description of the Effect*, ACS ISR 2004-12, by M. Giavalisco

11. Documentation and Web Site Updates

Several updates to ACS documentation should make it easier for observers to find information about ACS. First, the *ACS Instrument Handbook* (Gonzaga et al. 2005) has been reduced in length by about 15% without loss of content by consolidating information and moving some of the ACS calibration plan material to the ACS web site. The index for the handbook has also been updated, and a series of summary tables at the front of the handbook have been added to make it easier to find basic information about the instrument and its supported modes of operation. The *ACS Instrument Handbook* was released with the Cycle 15 *Call for Proposals*.

A short update to the *ACS Data Handbook* is planned (Pavlovsky et al. 2006). The update will include minor corrections to wording and syntax of some of the examples, as well as clarification of some of the text related to prism/grism reductions. The updated handbook is expected to be released in early 2006.

ACS long term usage statistics are now being charted on the STScI metrics web page at *http://www.stsci.edu/hst/metrics/SiUsage/ACS_LT*. This page contains both graphical and tabular information about requested ACS modes, filters, and exposure times.

Acknowledgments. Calibration of the ACS is truly a team effort. I thank the members of the ACS Calibration Team for their dedicated efforts to specify and conduct the ACS calibration program described in this article. I thank the PIs of the calibration programs listed in Tables 2-4 for their descriptions of the calibration programs. I am particularly grateful to Ron Gilliland, the ACS calibration lead, for his willingness to distill a large amount of calibration information and organize it into tabular form.

References

Gonzaga, S., et al. 2005, *ACS Instrument Handbook*, Version 6.0, (Baltimore: STScI)

Pavlovsky, C., et al. 2006, *ACS Data Handbook*, Version 5.0, (Baltimore: STScI)

Sirianni, M., et al. 2005, PASP, 117, 1049

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Empirical PSFs and Distortion in the WFC Camera

Jay Anderson

Rice University, Physics and Astronomy Department, MS-108, Houston, TX 77005

Abstract. This report summarizes two upcoming ISRs, one on empirical models of the WFC PSF and the other on the distortion solution. Before we can come up with an accurate distortion solution, we must first be able to measure accurate positions for stars. With a good PSF model, positions for well-exposed stars can be measured to about 0.01 pixel in a single exposure. To make the most of such accurate positions, we need a distortion solution which is at least as accurate. Distortion in the WFC comes from three primary sources: (1) the off-axis optics of the telescope, (2) a small perturbation introduced by each filter, and (3) a small irregularity in the CCD with a pattern that repeats every 68.270 columns. My solution addresses all three sources of distortion. I will describe how to access my solution and will also give advice on how best to use it. In addition, I will evaluate how breathing and secular variations introduce small additional distortions.

1. Introduction

Thanks to its excellent resolving power, *HST* can measure stellar positions that are about ten times more accurate than those that can be measured from the ground. In addition, it can do this for stars that are much fainter and in much denser fields than ground-based observatories can accurately observe. However, if we want to make full use of *HST*'s astrometric potential, we need a distortion solution that is at least as accurate as the star positions we hope to measure. Distortion solutions are also critical for applications other than stellar astrometry. Image stacking and mosaicking, modeling of extended sources, and source identification all require accurate distortion solutions.

2. PSFs

Astrometry is the main driver for an accurate distortion solution. While a solution accurate to 0.05 pixel is sufficient for most applications (drizzle, etc), if we cannot remove the distortion more accurately than we can measure the stars, then the position accuracy will be wasted. At the same time, we cannot measure an accurate distortion solution without accurate stellar positions. So, our first task in modeling the distortion is to measure accurate positions for stars. For this we will need to construct an accurate point-spread function (PSF).

The material in this section comes largely from an ISR that is nearly ready for publication about empirical models of the WFC PSF. This ISR will give information on how accurate our PSFs are, and how to download and use them.

2.1. A word about images

Please note that in the following PSF discussion, all analysis is done on the `_flt` images produced by the ACS pipeline. Stars in the `_drz` images cannot be well fit with PSF modeling, since each star image has been resampled in a different way by the drizzling

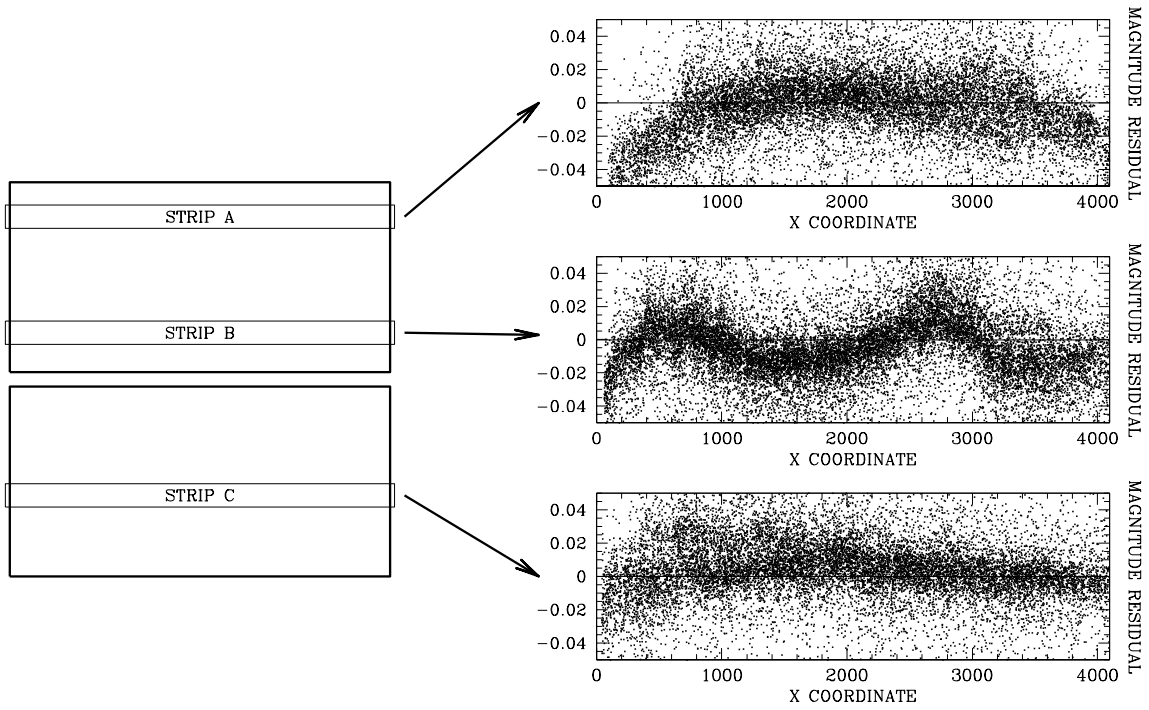

Figure 1: This shows the systematic photometric residuals that result from using a PSF that is constant across the chip to measure fluxes for stars. We show the residuals along three horizontal strips across the detector. The typical errors are ∼0.02 magnitude, but they can be as large as 0.04.

process. Although the `_flt` images are in a distorted frame, the stellar images can still be analyzed accurately by PSF fitting.

2.2. The need for better PSFs

In Anderson & King (2000, AK2000), we came up with a simple empirical model for the WFPC2 PSF. This model could be used to fit stars and measure positions with a systematic accuracy of 0.02 pixel. When we constructed a similar model for the ACS/WFC, we found that it could measure positions even better than this.

Our initial model for the WFC was just a single PSF which did not change shape with position on the detector (Anderson 2002). This spatially constant PSF model seemed to do well at first, but we did notice some spatial variation in the photometry which we initially thought was due to errors in the flat fields (see Figure 1). We later learned from Krist (2003) that variations in chip thickness cause variations in the amount of charge diffusion across the chip. This variation in turn perturbed the PSF, so that a single PSF model would not be valid everywhere. The core of the PSF can change sharpness significantly over spatial scales of 500 pixels. Photometry based on PSF fitting is particularly sensitive to the core, and such variations in core intensity can translate into systematic errors in photometry.

Since astrometry relies on core-shape even more than photometry does, we realized that we needed to construct more accurate PSF models. In AK2000, we constructed an array of 3×3 fiducial PSFs to treat the variation of the PSF across each 800×800 WFPC2 chip. The Krist ISR indicates that a much more extensive array will be necessary to deal with the spatial variations of the PSF in the WFC.

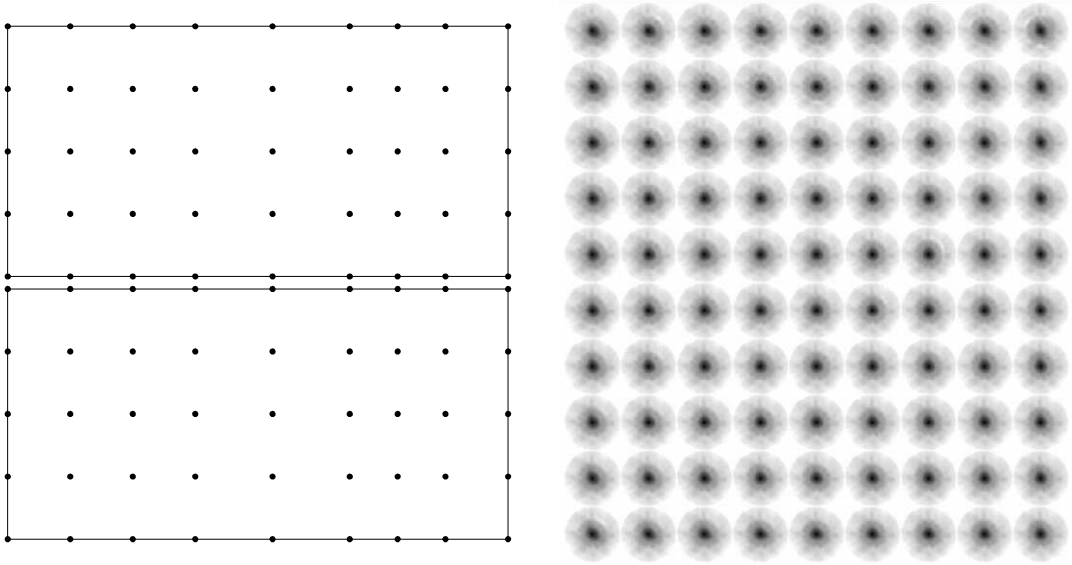

Figure 2: Left: the locations in the WFC chips where our fiducial models are constructed. Note that with models placed at the edges and corners we never have to extrapolate. Right: the `fits` file that contains our 9×10 array of PSF models.

2.3. The PSF models

We constructed an array of 9×5 fiducial PSFs to cover each chip. The fiducial PSFs are in general separated by 512 pixels, though we did move one of the fiducial points to coincide with a location on the chip where the thickness has an extremum (at $x = 2800$). Figure 2 shows the locations of the fiducial models and the array of PSFs that correspond to them. Note that the PSF is not assumed to be continuous across the gap.

Each fiducial PSF is represented empirically by tabulating the value of the PSF at a super-sampled (×4) array of grid points. The empirical PSFs we deal with are different from instrumental PSFs or Tiny Tim PSFs, which may require integration over pixels or convolution with diffusion kernels. Our "effective" PSFs include all of the pixel-integration effects implicitly and simply tell us what fraction of light should be recorded by a pixel that is centered at an offset $(\Delta x, \Delta y)$ from the center of a point source.

This model successfully represents the PSF from the core out to a radius of 12.5 pixels. We used a large dithered data set of a moderately dense globular cluster field to construct our PSFs, and arrived at the best model for the PSF in each region by optimizing the PSF to fit the images of all the stars found in that region of the detector.

On the right side of Figure 2, it is clear that to first order, the PSFs look quite similar. Figure 3 shows the average PSF for the F606W filter. On average, about 22% of a star's light will fall within the central pixel when the star is centered on that pixel.

In Figure 4, we subtract the average PSF from each of the fiducial models, to show how the PSF varies across the chip. In general, it varies in two ways. We see that the intensity of the core varies over a relatively short spatial scale, tracing the variation in charge diffusion properties. The map of the core intensity closely resembles the maps of the flat-field's overall intensity, which traces chip thickness and charge diffusion. The core is sharpest along $x = 2800$ where a star centered on a pixel will have 26% of its light in its central pixel. At around $(1500, 3000)$ the PSF suffers the most from diffusion and this fraction goes down to below 19%. This is roughly a ±10% variation and can have a significant effect on photometry if not included in the PSF model.

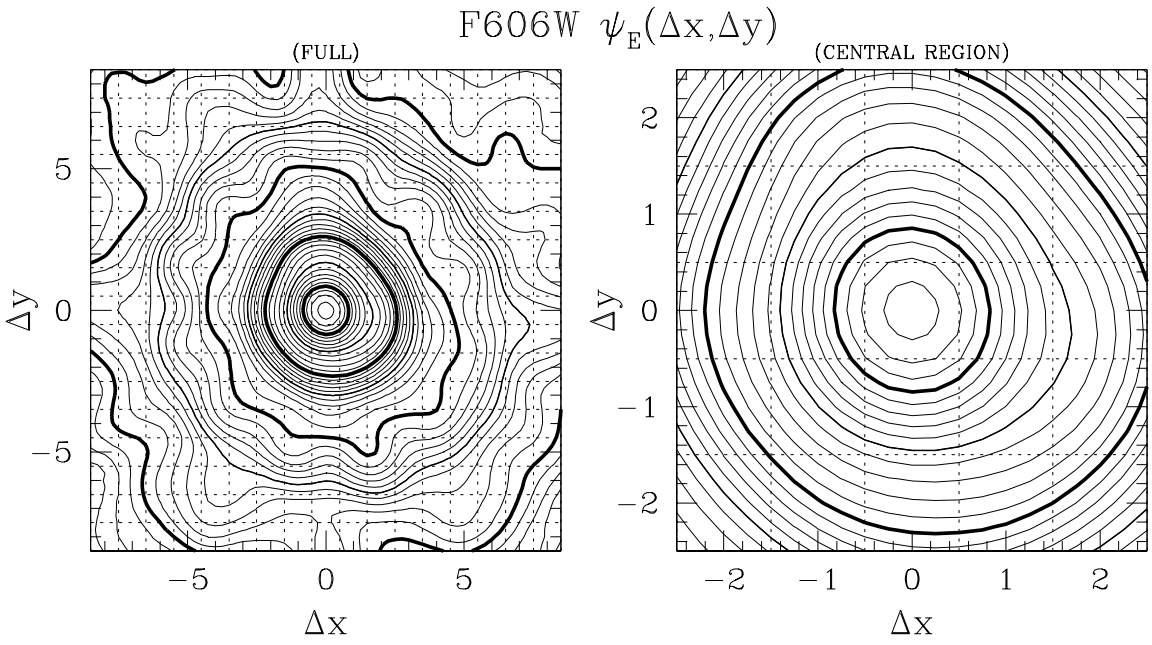

Figure 3: This shows a contour plot of the average F606W PSF. The heavy contours correspond to PSF intensities of 0.1, 0.01, 0.001, and 0.0001. The lightest contours correspond to 0.1 dex. Left: the full PSF over 8×8 pixels, Right: the central region. The dotted lines have a spacing of one pixel unit.

Outside of the core, we see the effects of geometric optics gradually changing the shape of the inner halo of the PSF (the region outside the core) from one corner of the chip to another.

2.4. Temporal variations of the PSF

We have found that this array of PSF models does a very good job dealing with the spatial variations in the detector and optics for the data set for which it was constructed. In the upcoming ISR, we will examine how stable the PSF is over time. We will show that the PSF can change over the short term because of breathing and over the long term because of secular focus variations.

The breathing variations can change the fraction of light in the core by about 5%: exposures in the beginning of an orbit can have sharper or broader cores than exposures at the end of the orbit. The longer-term focus variations can introduce a 10% variation in the PSF over a period of months to years. These variations could introduce photometric errors of $\sim$3% if not addressed.

We examine these changes in detail and find that although the fraction of light in the core changes by quite a bit, the rest of the PSF structure is still quite similar. This is somewhat to be expected since much of the PSF variation across the chip comes from the diffusion characteristics of the detector, which do not change with breathing or focus. In the PSF ISR we will present a heuristic way of dealing with the temporal variability.

2.5. Products available soon

The upcoming ISR will explain our PSFs in more detail and will also provide links to the empirical PSFs we have constructed for several WFC filters: F435W, F475W, F606W, F625W, F658N, and F814W. If there exist well-dithered data sets of a relatively dense star field through other filters, it would not be hard to expand this set to cover other filters. We

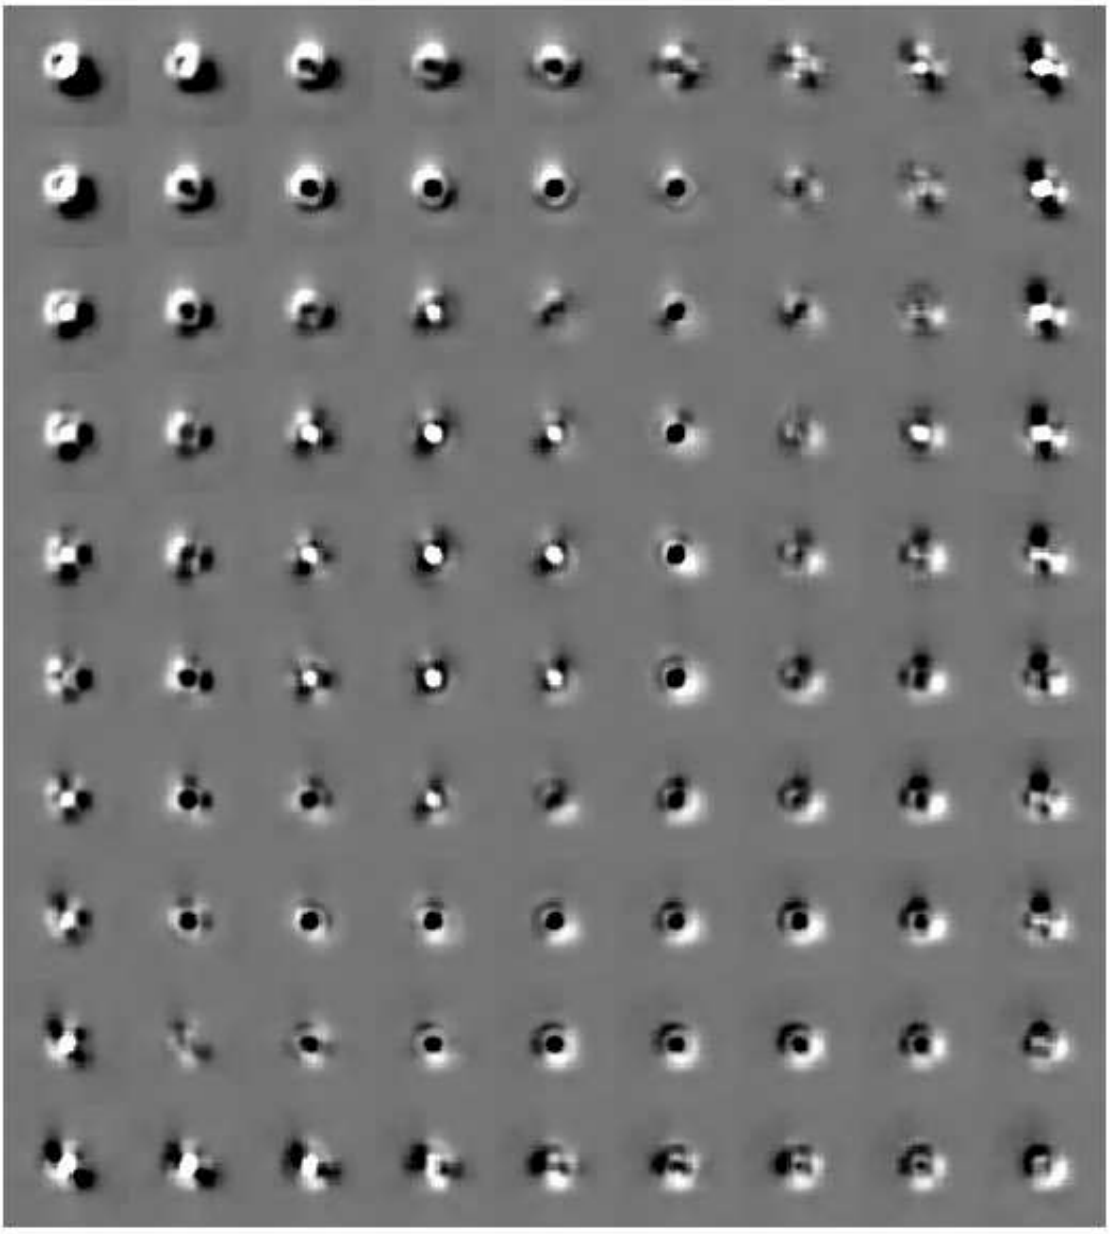

Figure 4: This shows how each of the 9×10 PSFs differs from the average. Here we show only the inner 6 pixels (in radius) of the PSF. Dark here implies more flux.

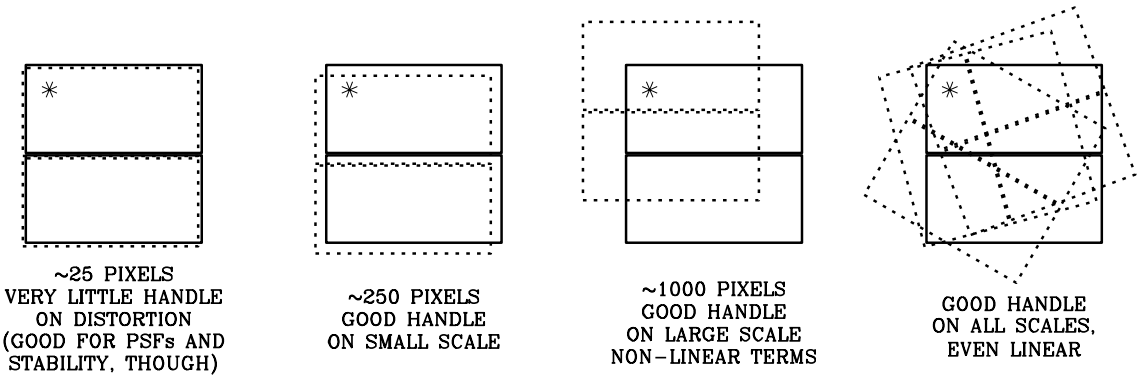

Figure 5: This plot shows different kinds of data sets that can be used to probe distortion at different scales. The idea is that the same star is observed in different places in the detector in different exposures, allowing us to probe the distortion. The fields observed have many thousands of stars.

have not yet compared our PSFs to the predictions of Tiny Tim models, but that would be interesting to do.

The ISR will also provide software that uses these PSFs to measure positions and fluxes for stars in WFC images. This routines can measure reasonably bright stars (total fluxes over 10^4 e^-) with positional errors less than 0.01 pixel in each coordinate and with photometry errors less than 0.01 magnitude. These errors represent the error in a single exposure. By combining multiple exposures one can reduce this even more. The PSFs can also be used for detailed analysis of resolved objects or deconvolution.

3. The Distortion Solution

As mentioned above, astrometry is the main driver for an extremely accurate distortion solution. But other applications will also benefit from the best possible solution: image rectification, stacking of multiple exposures to go deep, source identification, and lensing-type analyses. These other applications may not require 0.01-pixel precision, but errors of 0.05 pixel could introduce unnecessary problems.

3.1. The complications of self-calibration

The easiest way to calibrate the distortion in an instrument would be to compare an observation against an absolute standard field where all the star positions are known in a calibrated frame to a high accuracy. Since no such fields exist for *HST*, we must rely on self-calibration to tease out the distortion solution.

While comparison against an absolute standard can reveal distortions on all scales, self-calibration often requires different data sets to examine different aspects of the distortion. Figure 5 shows different kinds of data sets that can be used to solve for different aspects of the distortion solution.

3.2. Observing programs to pin down the solution

Soon after ACS was installed, *HST* calibration program 9028 imaged the center of 47 Tuc through the F475W filter. Meurer et al. (2002) used these observations to construct the first polynomial-based solution for the distortion. They were able to remove most of the distortion with a fourth-order polynomial, but they did note that there were still some higher-order systematic residuals at the 0.05-pixel level.

GO program 9443 was taken three months later in order to pin down the linear terms of the solution and to examine the solution's behavior with wavelength. Anderson (2002) saw the same pattern of residuals, but noticed that this pattern was different for observations through different filters.

GO program 10424 was taken in December 2004 in order to construct empirical library PSFs and to improve the distortion solution for the most used filters, F606W and F814W. It turned out that the observations were taken right before the *HST* focus was re-adjusted, so the data are not very useful for probing the average properties of the WFC. But the observations can tell us what extreme conditions can do to PSFs and the distortion solution.

3.3. The current best solution

This write-up is too short to go into how distortion solutions are constructed from the various kinds of data sets in Figure 5. So, I will just summarize the properties of the final solution:

- The backbone of the solution is a fourth-order polynomial. Since the x and y axes of the WFC are off from perpendicular by 9°, the largest terms are linear. But there is still considerable power in the quadratic and other high-order terms. For instance, the amplitude of one of the quadratic terms is 35 pixels.

- In addition to the global polynomial solution, we found that each filter introduces a perturbation on this solution. One can think of each filter having a distortion fingerprint. The amplitude of this perturbation is typically 0.05 pixel and it is spatially coherent over 150 pixels or so. Each filter also has a slight shift of up to 0.5 pixel and a slightly different scale, with FOVs differing by up to 0.5 pixel.

- A final contributor to the distortion is a pattern in the detector itself. The pixels were drawn onto the chips with a stencil that had slight irregularities in the pixel sizes. These irregularities are less than 0.01 pixel in amplitude (typically 0.001 pixel), but they are still worth removing. The pattern repeats every 68.270 columns, and we provide a simple astrometric correction that removes it. The flat fields should remove the pattern's photometric signature.

3.4. Accessing the solution

The distortion solution is provided in the form of a `FORTRAN` subroutine. It can be found at: `http://spacibm.rice.edu/~jay/WFC/FORTRAN/DISTORTION.`

The routine has six arguments: the raw x and y positions, the chip number i and the filter number k. It returns the distortion-corrected positions. The directory also contains a routine that performs the inverse distortion correction.

If `FORTRAN` is not your preferred language, I have also generated some simple `fits` images that can serve as a simple look-up table for the distortion correction, so that the solution can be accessed within any programming environment. These images are not on the website, but feel free to e-mail about them.

3.5. Accuracy of the solution

The solution was initially constructed for the F475W filter, but other filters have been added. The solution for F606W and F814W has recently been improved, and the solutions for F435W, F625W, and F658N are quite good. The solutions for F555W, F850L, F775W, F660W, F502N, and F550M are based on only a few images and are not as accurate.

The solutions for the best six filters are in general accurate to better than 0.01 pixel, apart from breathing and other issues. Breathing affects the linear and other low-order terms, but it does not seem to affect the filter-specific part of the distortion. The filter fingerprint seems to be quite stable over time.

3.6. Time variation of the distortion

The distortion solution is usually constructed from a single data set at a given point in time. Below we will look at how much the solution varies over time.

Short-term variations GO program 10424 (PI Richer) took 378 exposures of a moderately dense field in globular cluster NGC 6397 over the period of a month. This data set is similar to the leftmost panel of Figure 5. It cannot be used to construct the distortion solution, but it can tell us a lot about how stable the solution is over time. To study this, we measured all the stars in all the images and constructed an average position for each star. We then looked at where each star fell in each image compared to the average, in order to see if there were any astrometric variations that could be traced to the distortion solution.

The linear terms show the biggest variation. As *HST* moves towards or away from the target, the FOV expands or contracts by as much as 0.05 pixel because of velocity aberration. This can be removed using the `VAFACTOR` keyword in the image header. Aside from this, the field changes by less than 0.01 pixel during the month of observations, meaning the linear terms are stable to 5 parts in 10^6!

We also looked at the quadratic terms and noticed that the solution seemed to vary by about 0.02 pixel. These changes were correlated with changes in the focus, so it is hard to say if the distortion solution is changing, or if the PSF is changing. At some level, this is just a question of semantics. The good news is that the solution appears to be extremely stable.

Long-term variations There is a field in 47 Tuc that has been observed well over 350 times by the WFC (over 160 times through F606W), at a variety of orientations and offsets. We gathered together all of these observations in order to study how the distortion solution may have changed over the 3.5 years that the WFC has been on board *HST*.

Figure 6 shows several diagnostic quantities derived by comparing each image to all the others. In addition to the astrometric diagnostics, in the top panel we also show how the focus has changed over the observation period. The second panel shows the raw linear scale of each image compared to the average. The clear annual trend results from velocity aberration, and is proportional to the dot-product of the Earth's motion with the direction of 47 Tuc. The amplitude for this effect is about 0.12 pixel.

If we remove the expected trend (using the `VAFACTOR` keyword), we see in the middle panel that there is still some unexplained variation of about 0.03 pixel in the FOV radius. We see the same irregular 0.03 pixel variation in the on-axis linear skew (which corresponds to an expansion along the x axis at the expense of the y axis).

The bottom panel shows the off-axis skew, which corresponds to a linear shear. This quantity exhibits a clear secular trend, with a slope of about 0.03 pixel per year. Over the lifetime of the WFC, this term has changed by over 0.1 pixel. Since the original solution was constructed from data taken in mid 2002, this off-axis linear term of the solution is currently in error by about 0.1 pixel. This was found to be a problem in the construction of the UDF a year ago. In my upcoming ISR on the distortion solution, we will provide a correction to the solution for this.

Stability with focus The observations in GO program 10252 were taken of a dense field in Omega Cen in order to pin down the distortion solution for F814W and F606W. As it happened, the detector was very much out of focus at the time and the observations were not good for the intended purpose. But they are still useful for showing us what kind of errors we should expect in conditions of extreme focus.

We will show in the upcoming ISR that the solution is largely the same, but there are some low-spatial-frequency errors of 0.01 pixel which can reach 0.02 pixel at the edge of the chip. As we saw above, this variation could easily be related to variations in the PSF, which we would expect when the focus changes.

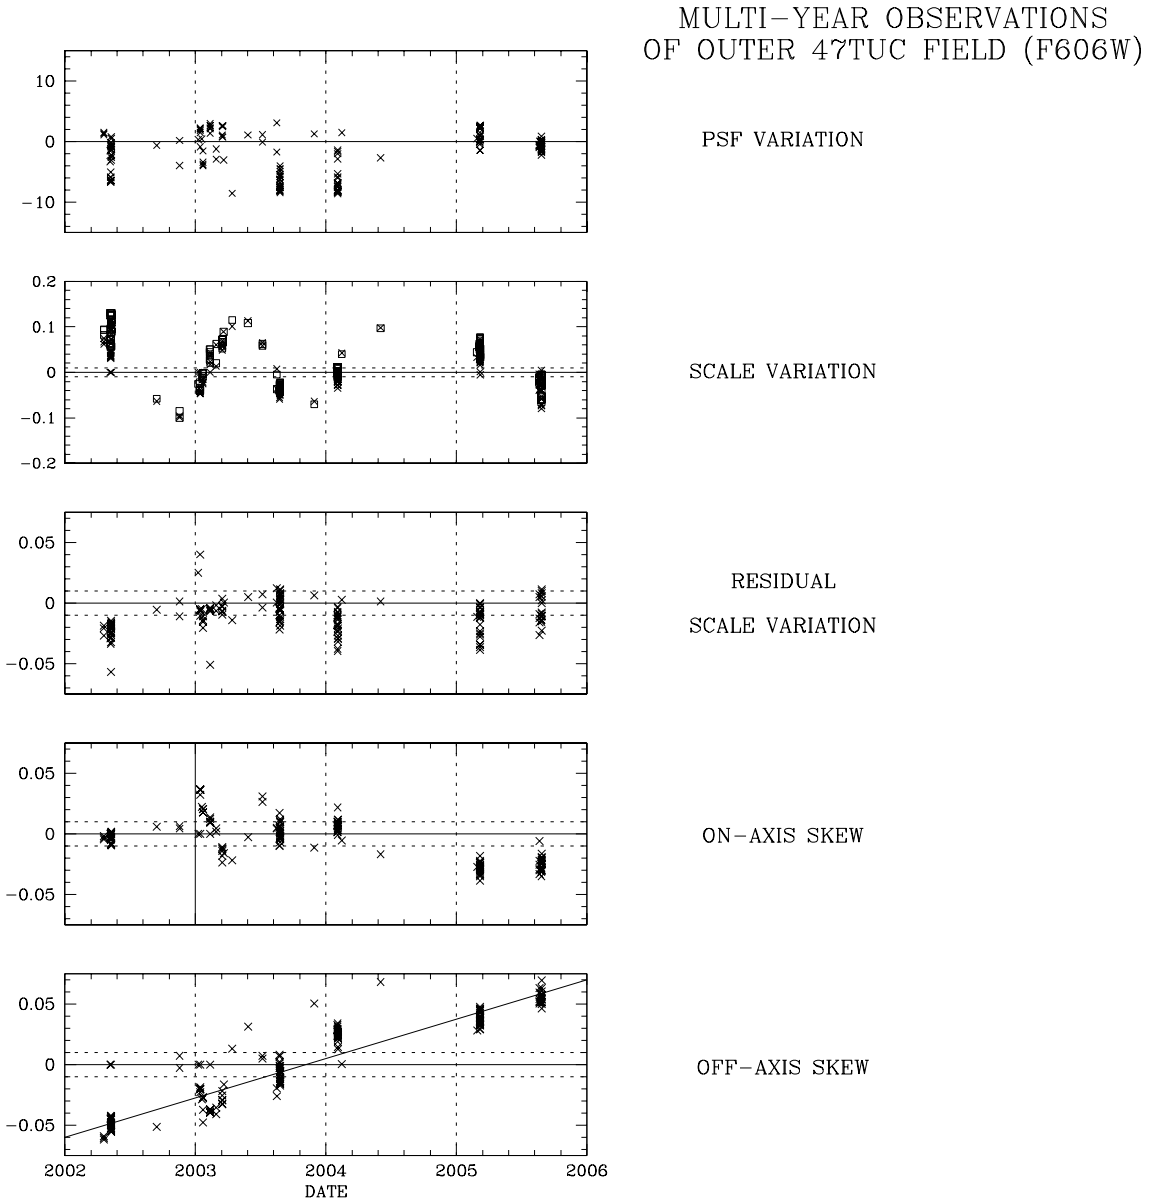

Figure 6: This plot shows how terms in the astrometric solution have changed over the 3.5 years that the WFC has been on board *HST*. The top plot shows the focus (in terms of the intensity of light in the core). The second plot shows the total variation of the plate scale over time (in terms of expansion of the FOV). The middle plot shows the same variation, with the `VAFACTOR` contribution removed. The fourth plot shows the on-axis linear skew, and the bottom plot shows the off-axis linear skew.

Again, the good news is that the distortion solution does not appear to change very much at all. It seems to be globally stable to better than 0.02 pixel, if you allow for some variation of the linear terms.

3.7. Using the distortion correction

The distortion correction presented here is available at the website given in §3.4 in the form of a `FORTRAN` program. It has been included in the drizzle stage of the ACS pipeline for quite some time now.

Astrometry is the main application that will require users to implement the distortion correction themselves. Measuring positions in the `_drz` images will typically result in errors of ∼0.05 pixel, on account of the resampling that has been done to the stellar images. Using a good PSF on the `_flt` images can yield systematically accurate positions to better than 0.01 pixel. The only complication is that the positions have to be corrected for distortion. The upcoming PSF ISR will provide our empirical PSFs and a program to measure `_flt` images with them. The program will even perform the distortion correction.

When doing astrometry with the WFC it is important to keep in mind the limitations of the distortion correction. As we have shown, the distortion correction for the well calibrated filters is globally accurate to better than 0.02 pixels, if allowance is made for variation of the linear terms. We recommend using 6-parameter linear transformations when comparing positions in one distortion-corrected frame with those in another. This should only be a hardship for programs with fewer than 3 reference stars.

Even though the global solution appears to be good to 0.02 pixel, it is often still preferable to perform astrometric measurements as locally as possible. The largest systematic errors will result when comparing the position of a star in one corner of a chip with that of a star in another corner. If you can keep your differential astrometric measurements more local (say, over 200 pixels or so), then any systematic errors in the solution will be minimized.

The upcoming ISR on distortion will go into more depth on the stability of the solution. It will also recommend observation strategies for programs with various astrometric goals.

References

Anderson, J., & King, I. R., 2000, PASP, 112, 1360

Anderson, J., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 13

Anderson, J., & King, I. R., 2003, PASP, 115, 113

Anderson, J., & King, I. R., 2004, *Instrument Science Report* ACS 2004-15 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Krist, J., 2002, *Instrument Science Report* ACS 2002-06 (Baltimore: STScI)

Meurer, G., Lindler, D., Blakeslee, J. P., Cox, C., Martel, A. R., Tran, H. D., Bouwens, R. J., Ford, H. C., Clampin, M., Hartig, G. F., Sirianni, M., & De Marchi, G., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 65

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Modeling and Correcting the Time-Dependent ACS PSF

Jason Rhodes

Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109

Richard Massey, Justin Albert, James E. Taylor

California Institute of Technology, 1200 East California Blvd, Pasadena, CA 91125

Anton M. Koekemoer

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD, 21218

Alexie Leauthaud

Laboratoire d'Astrophysique de Marseille, BP 8, Traverse du Siphon, 13376 Marseille Cedex 12, France

Abstract. The ability to accurately measure the shapes of faint objects in images taken with the *Advanced Camera for Surveys* (ACS) on the *Hubble Space Telescope* (HST) depends upon detailed knowledge of the Point Spread Function (PSF). We show that thermal fluctuations cause the PSF of the ACS Wide Field Camera (WFC) to vary over time. We describe a modified version of the *TinyTim* PSF modeling software to create artificial grids of stars across the ACS field of view at a range of telescope focus values. These models closely resemble the stars in real ACS images. Using ~ 10 bright stars in a real image, we have been able to measure HST's apparent focus at the time of the exposure. *TinyTim* can then be used to model the PSF at any position on the ACS field of view. This obviates the need for images of dense stellar fields at different focus values, or interpolation between the few observed stars. We show that residual differences between our *TinyTim* models and real data are likely due to the effects of Charge Transfer Efficiency (CTE) degradation. Furthermore, we discuss stochastic noise that is added to the shape of point sources when distortion is removed, and we present *MultiDrizzle* parameters that are optimal for weak lensing science. Specifically, we find that reducing the *MultiDrizzle* output pixel scale and choosing a Gaussian kernel significantly stabilizes the resulting PSF after image combination, while still eliminating cosmic rays/bad pixels, and correcting the large geometric distortion in the ACS. We discuss future plans, which include more detailed study of the effects of CTE degradation on object shapes and releasing our *TinyTim* models to the astronomical community.

1. Introduction and motivation

Accurate shape measurements of faint, small galaxies are crucial for certain applications, most notably the measurement of weak gravitational lensing. Quantifying the slight distortion of background galaxies by foreground matter relies on detecting small but coherent changes in the shapes of many galaxies (see Refregier 2003 for a recent review). To extract the lensing signal, it is crucial to remove instrumental effects from galaxies' measured shapes. On the *Hubble Space Telescope* (HST), these include:

- Convolution of an image with the telescope's Point Spread Function (PSF).
- Geometric distortion of an image. This is particularly large in the *Advanced Camera for Surveys* (ACS) because of its location off HST's optical axis.
- Trailing of faint objects in the CCD readout direction due to degraded Charge Transfer Efficiency (CTE).

In this proceeding, we describe a method to model and correct for the telescope's temporally and spatially varying PSF. The geometric distortion has already been shown to be successfully removed during image processing by *MultiDrizzle* (Koekemoer et al. 2002). Removing the distortion does change the PSF, and we present recommendations to minimize stochastic changes introduced during the repixellization stage of image processing. The effect of continuing CTE degradation on galaxy shapes is only becoming apparent as the ACS spends longer in orbit, and is not yet completely understood. That is therefore beyond the scope of this proceeding. A separate method to remove CTE effects will be presented in Rhodes et al. (2006), and the application of all these corrections in a weak lensing analysis will be presented in Massey et al. (2006). Other branches of astronomy, including stellar photometry in crowded fields, the study of AGN, and proper motions also require detailed knowledge of the PSF and will benefit from the models we describe here.

In weak lensing, to deconvolve galaxy shapes from the PSF, we must accurately know the shape of the PSF at the position of the object and at the time of the observation For example, see Rhodes, Refregier & Groth (2000) for a description of the method we use on the Cosmic Evolution Survey (COSMOS; Scoville et al. 2006) images we use to test the PSF models we describe in this paper. If the HST PSF were stable over time, it would be straightforward to build a catalog of stellar images across the entire field of view. However, thermal fluctuations in HST that change its effective focus (the distance between the primary and secondary mirrors) lead to temporal PSF variations. As an example, Figure 1 shows the PSF pattern in two sets of COSMOS images. The left hand panel shows stars from images taken when the telescope was near optimal focus, and the right hand panel shows stars observed when the telescope was several microns below optimal focus. Each tick mark in the figure represents the ellipticity of one star, measured using the standard weak lensing definition,

$$|e| = \frac{\left[(I_{xx} - I_{yy})^2 + (2I_{xy})^2\right]^{1/2}}{I_{xx} + I_{yy}} \ , \tag{1}$$

where star's weighted second order moments

$$I_{ij} = \frac{\sum wIx_ix_j}{\sum wI} \tag{2}$$

involve summations over all pixels. I is the intensity of a pixel, w is some weighting function (in our case a Gaussian with a width of about the FWHM of the PSF), and x_i is the distance of a pixel from the centroid of the object. It is apparent from Figure 1 that changes in the PSF over time are sufficient for a temporally stable ACS PSF model to be inadequate in demanding applications, when using data collected over a period of more than a few days. Other effects, including CTE, introduce additional variation on longer time scales.

This proceeding is organized as follows. In §2. we introduce the *TinyTim* software package that we used for PSF modeling, and discuss modifications that we have made. In §3. we discuss *MultiDrizzle* and how to minimize the aliasing of point sources that occurs during distortion removal. In §4. we show how we have used our *TinyTim* models to quantify the temporal variation of the ACS PSF, and describe how the same models can be used to correct for it. In §5. we draw conclusions and outline a plan to release our *TinyTim* PSF models.

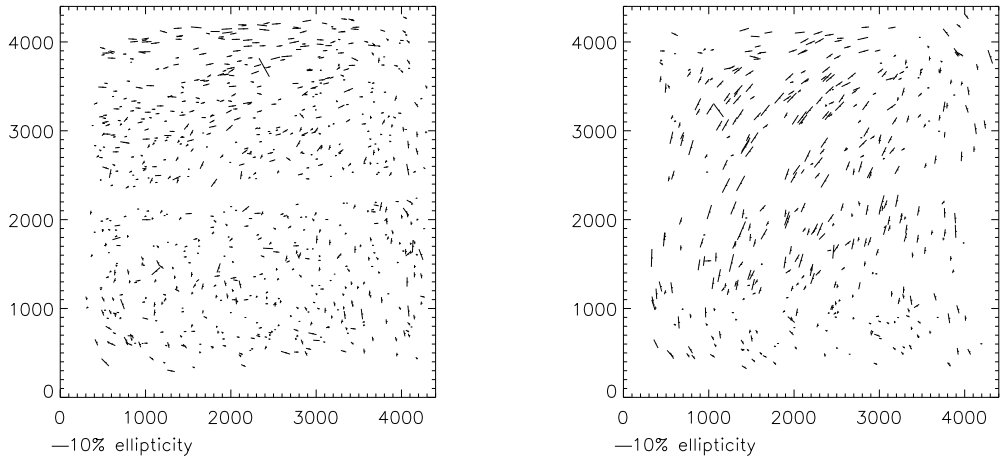

Figure 1: The ellipticity of stars in the COSMOS survey observed with the ACS WFC while HST happened to be at nominal (left panel) and low (right panel) focus. The orientation and size of each tick mark represents the ellipticity of one star; both panels contain stars from several different fields. The difference in the PSF patterns is apparent, and demonstrates the need for a time dependent PSF model.

2. TinyTim PSF models

We have adapted version 6.3 of the *TinyTim* software package (Krist & Hook 2004) to create simulated images of stars. *TinyTim* creates FITS images containing one or more stars that include the effects of diffraction, geometric distortion, and charge diffusion within the CCDs. By default the images appear as they would in raw ACS data; they are highly distorted and have a pixel scale of 0.05 arcseconds per pixel. We have written an IDL wrapper to undo the distortion, resample the images, and combine adjacent PSFs to mimic the effects of dithering. The wrapper can also run *TinyTim* multiple times and create a grid of PSF models across the whole ACS field of view. We insert our artificial stars into blank images with the same dimensions and FITS structure as real ACS data, thereby manufacturing arbitrarily dense starfields.

This basic pipeline calculates a diffraction pattern (spot diagram), distorts it, and adds charge diffusion; all three effects usually depend upon the position of the star in the ACS field of view. We have made two versions of artificial starfields with important changes to this basic pipeline. The deviations from this default are:

- In order to examine the effects of the distortion removal process, we have created a version of our *TinyTim* starfields where each star has an identical diffraction pattern and charge diffusion, but a geometric distortion determined by the location of the PSF within the ACS field of view. Once the geometric distortion is removed (for example by running the field through *MultiDrizzle*), these stars should all appear identical.

- In order to correct data, we have created a second version of our *TinyTim* starfields that do not contain the effects of geometric distortion at all, instead modeling stars as they would appear after a perfect removal of geometric distortion. Conversion between distorted and non-distorted frames, which is necessary to simulate charge diffusion in the raw CCD, was performed using very highly oversampled images. This avoids stochastic aliasing of the PSF (see §3.), and minimizes noise in the PSF models.

We describe the application of these simulations in the following two sections.

3. Optimization of MultiDrizzle parameters for Weak Lensing Science

MultiDrizzle is used to combine dithered exposures, remove cosmic rays and bad pixels, and eliminate the large geometric distortion in ACS WFC images (Koekemoer et al. 2002). However, the transformation of pixels from the distorted input image to the undistorted output plane can introduce significant "aliasing" of pixels if the output pixel scale is comparable to the input scale. When transforming a single input image to the output plane, point sources can be enlarged, and their ellipticities changed by several percent, depending upon their sub-pixel position. This is one of the fundamental reasons why dithering is recommended for observations, since the source is at a different sub-pixel position in different exposures, hence the effects are mitigated to some extent when all the exposures are combined. However, the remaining effect in combined images is still quite sufficient to prevent the measurement of small, faint galaxies at the precision required for weak lensing analysis.

Of course, such pixellization effects are unavoidable during the initial exposure, when the detector discretely samples an image. However, it is clearly desirable to minimize related effects during data reduction. The effect on each individual object depends on how the input and output pixel grids line up. Indeed, this can be mitigated by using a finer grid of output pixels (*e.g.* Lombardi et al. 2005). The reduction in pixel scale (which will cause a corresponding increase in computer overheads) can be performed in conjunction with simultaneously "shrinking" the area of the input pixels that contains the signal, by making use of the *MultiDrizzle* `pixfrac` parameter and an appropriate kernel.

We have run a series of tests on the simulated PSF grids described in §2. to determine the optimal values of the *MultiDrizzle* parameters specifically for weak lensing science. To this end, we first produced a grid of stars that ought to look identical after the removal of geometric distortion. Figure 2 shows the "aliasing" produced when the distortion is removed from a single image. We have also created a series of four dithered input images (with the linear dither pattern used for the COSMOS survey); the scatter in the ellipticities of the output stars is then approximately half as big. This confirms the idea that the repixellization adds stochastic noise to the ellipticity when the four sub-pixel positions are uncorrelated. For weak lensing purposes, this noise is still substantial. With enough dithered input images, the scatter of ellipticities could be reduced further, but this is not feasible for most observations.

We then ran a series of tests using *MultiDrizzle* on the same input image but with a range of output pixel scales, convolution kernels, and values of `pixfrac`. We then measured the scatter in the ellipticity values in the output images. The smaller that scatter, the more accurately the PSF is represented. We found the results were not strongly dependent on the choice of`pixfrac` and settled on a value of 0.8 for that parameter. We show in Figure 3 that PSF stability is improved dramatically by reducing the output pixel scale from 0.05 arcseconds (the default) to 0.03 arcseconds. There is only a very slight gain in going to smaller output pixel sizes and the storage requirements rapidly become problematic. The gain in going to smaller pixel scales is more stable with a Gaussian kernel for than with the default square kernel. Therefore, for weak lensing work, we recommend an output pixel scale of 0.03 arcseconds, `pixfrac`=0.8, and a Gaussian kernel in order to best preserve the PSF during this stage of image reduction. We note that, despite its clear advantages for weak lensing studies, the Gaussian kernel does have some general drawbacks, such as the introduction of more correlated noise which may not be desirable for other types of science where minimization of correlated noise is important.

4. Quantification of the PSF and focus variability

We can measure HST's effective focus at the time of an observation using $\sim$ 10 fairly bright stars in the field. We first created dense grids of artificial stars across the ACS field

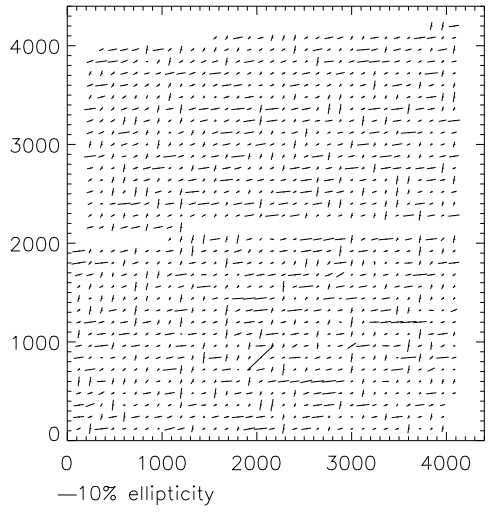

Figure 2: "Aliasing" of the PSF introduced when transforming a single distorted input image to the undistorted output frame. The tick marks represent the ellipticity of stars that have undergone identical diffraction in the telescope's optics and should therefore look identical. The only difference between stars is their sub-pixel position when their geometric distortion is removed and the images are combined. The apparent difference between the tick marks shows that the PSF is changed. The problem can be ameliorated by altering several of the *MultiDrizzle* settings, in particular reducing the output pixel scale.

of view. By changing the separation of the primary and secondary mirrors in *TinyTim*'s raytracing model, these models were made at successive displacements of the focus from nominal, from -10μm to $+5\mu$m in 1μm increments. These are reasonable bounds on the maximum extent of physical variations in HST. The stars are created without geometric distortion, to avoid the noise that would have been introduced had it been necessary to carry out geometric transformations on the stellar fields. We then compare the ~ 10 bright stars in each COSMOS field to the *TinyTim* PSF grids at each focus value. Minimizing the difference in ellipticity between the models and the data finds the best fit value of telescope focus for that particular field. Tests on observations of dense stellar fields that contain many suitable stars show that this procedure is repeatable with an rms error $\sim 1\mu$m, when using ten different stars repeatedly selected at random.

Figure 4 shows our estimation of HST's focus in microns away from nominal during several months in Cycles 12 and 13, using a uniform set of COSMOS images. HST was not manually refocussed during this time, but the apparent focus still oscillates. At times, the oscillations seem periodic, but there are also sharp jumps and more erratic behavior. The random component probably depends in a complex fashion upon the orientation of HST with respect to the sun and the Earth during preceding exposures, and we do not believe that it can be easily predicted in advance.

Note that the uncertainty on the focus value during any single pointing is quite large; more so than the tests on dense stellar fields would suggest. A major component of this error is undoubtedly the $\sim 3\mu$m thermal fluctuations in focus that HST experiences during each orbit due to "breathing". The COSMOS images are all taken with a total exposure time of one orbit, and the apparent focus therefore represents the integral of a gradually changing PSF. We have not been able to investigate focus changes on short time periods and, given this behavior, it is even more curious that long term patterns are so clearly present in Figure 4.

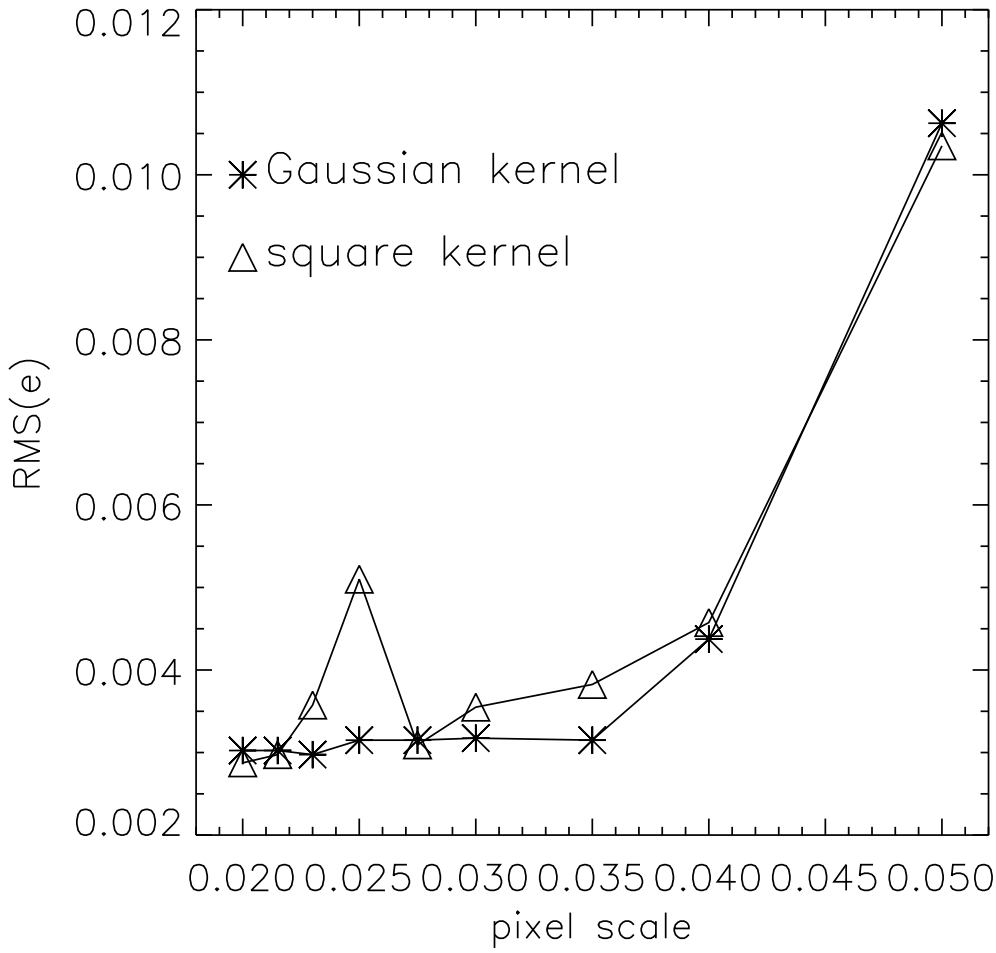

Figure 3: RMS ellipticity introduced during the process of removing geometric distortion and combining dithered images, for a range of *MultiDrizzle* parameters. Lower values show more stable behavior of the PSF during this process. Based on this plot, we recommend a Gaussian kernel and an output pixel size of 0.03 arcseconds in order to minimize the effect of undersampling on the PSF and produce images that are optimal for weak lensing science. We note that the Gaussian kernel introduces significant additional correlated noise relative to the square kernel, which is not important for weak lensing science but may not be optimal for other types of science.

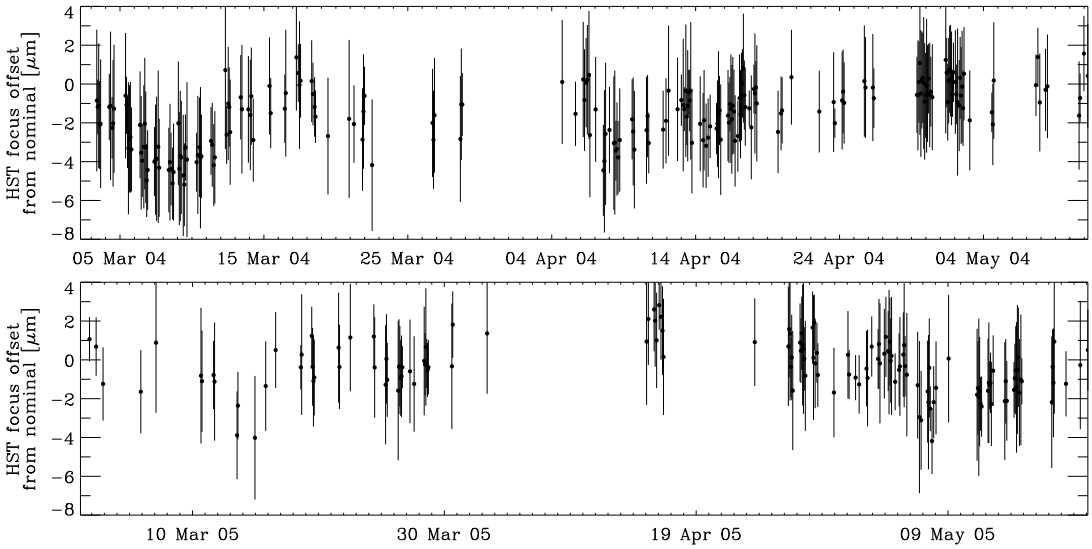

Figure 4: Apparent offset of HST from nominal focus during COSMOS observations in cycles 12 and 13. We describe the procedure by which we estimate the telescope's focus in §4.

Figure 5 shows the focus values determined for all of the COSMOS fields. The obvious clustering of focus values in that plot is due to the observing strategy used for COSMOS, in which data was typically taken in chunks of a few fields at a time. Adjacent fields are likely to have been taken at similar times, and therefore tend to have a similar focus values.

Figure 6 shows the *TinyTim* models at focus -3μm and all the stars from the COSMOS fields with a best-fit focus value of -3μm. The COSMOS stars have been averaged in a spatial grid of approximately 600 × 600 0.03 arcsecond pixels. There is good agreement between the models and the stars over most of the field. The agreement is not as good in the boxed area near the center of the field. We believe this is due to a degradation of the CTE of the ACS CCDs. This degradation causes trailing of low flux objects in the readout direction (the y direction). The effect is most pronounced the further away the object is from the readout registers at the bottom and top of the field (Mutchler & Sirianni 2005; Riess & Mack 2004). This causes the objects to be elongated vertically. Thus, fainter COSMOS stars appear more elongated in the y direction at the center of the field than the *TinyTim* models, which do not include the effects of CTE. Note that this does not affect the estimation of focus positions, because the bright stars matched to our PSF models are less affected by CTE than the faint sources. We are currently exploring ways to correct for CTE in all objects, and will publish the results in Rhodes et al. (2006).

5. Conclusions and Future Work

We have shown that *TinyTim* can produce model PSFs that are very close to those observed in real data (for example the COSMOS 2-Square Degree Survey). This required some modifications to the *TinyTim* code, most importantly adding the ability to mimic the distortion correction and dithering normally implemented via *MultiDrizzle*, and to produce grids of PSFs across the entire ACS WFC field of view. We used *TinyTim* model stars to find the best values for *MultiDrizzle* and found that using a Gaussian kernel, `pixfrac` = 0.8, and an output pixel scale of 0.03 arcseconds greatly reduced the "aliasing" of point sources introduced during repixellization.

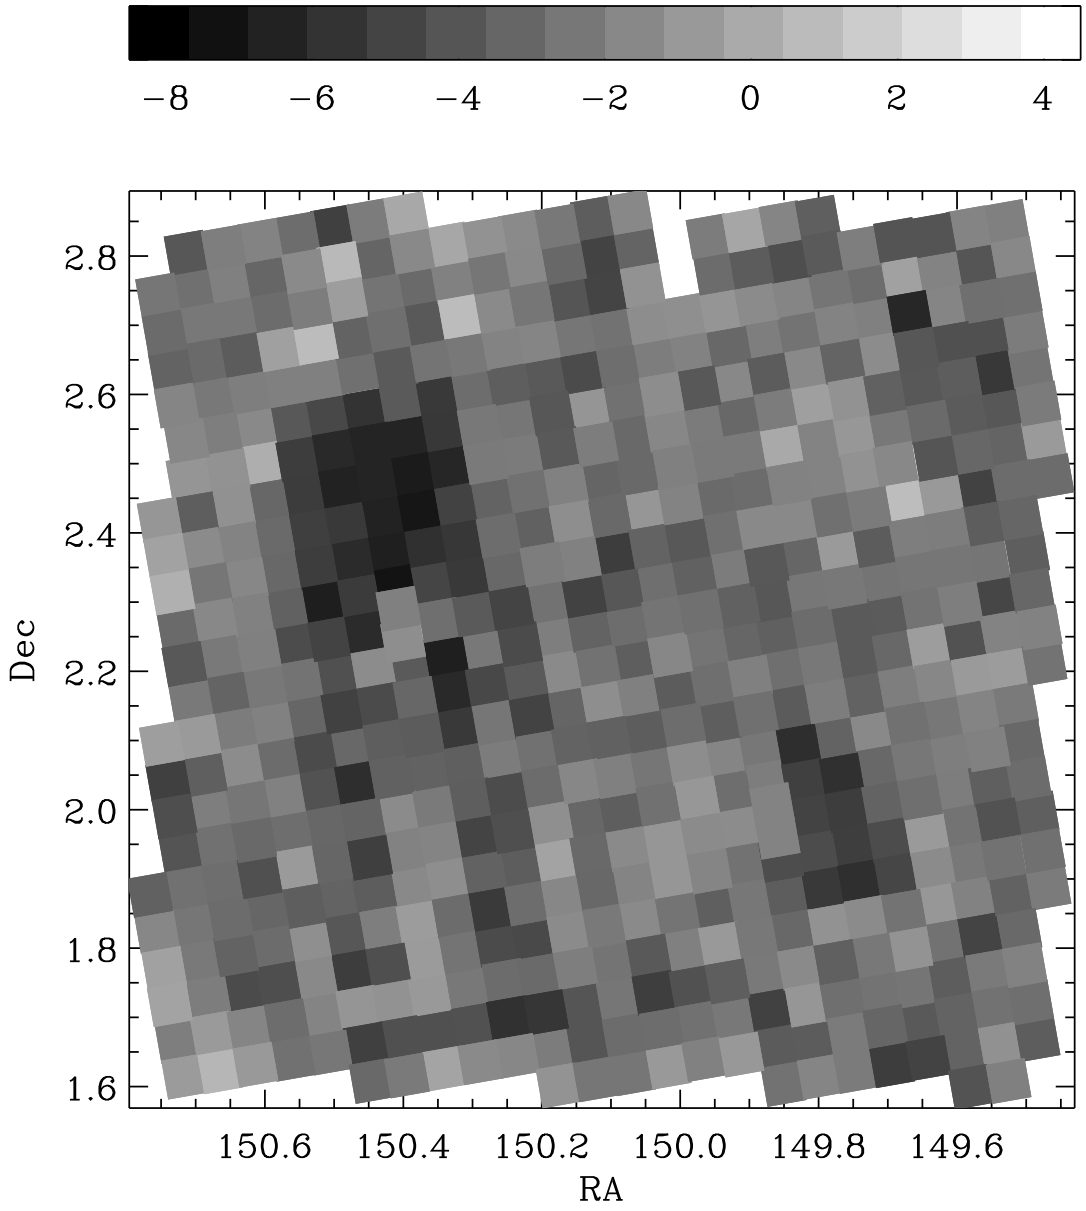

Figure 5: Apparent offset of HST from nominal focus in all of the COSMOS fields. COSMOS was taken in chunks, a few fields at a time. Therefore, the focus values cluster because fields taken close to each other in time tend to have similar focus values. Despite having only about 10 stars per COSMOS field which are suitable for measuring PSF, it is apparent from the clustering of focus values in this plot that we can make a decent estimate of the focus value for individual COSMOS fields using the models and techniques outlined in this paper.

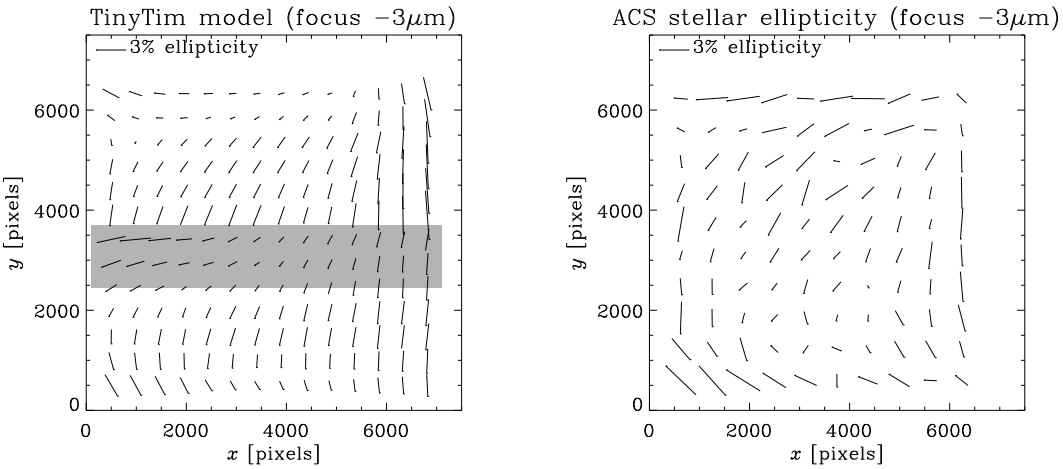

Figure 6: The *TinyTim* PSF models (left panel) for a focus value of -3μm and observed stars (right panel) in COSMOS fields with a similar apparent focus. There is good agreement between the data and the models over much of the ACS field. The shaded area near the center of the chip does not show good agreement and this is likely due to the effects of degradation of the Charge Transfer Efficiency (CTE) in the ACS WFC CCDs.

Discrepancies between our models and the COSMOS data can be attributed largely to a degradation in the ACS CTE since launch. We are currently studying this problem and will present a complete PSF solution including how to correct for CTE in Rhodes et al. (2006). We plan to correct science images for CTE on a pixel-by-pixel basis, like the Bristow code developed for STIS (Bristow et al./ 2002), moving charge back to where it belongs, rather than including the effects of CTE in our model PSF (like Rhodes et al. 2004). Thus, the model PSFs we present here are the ones we plan to use in our weak lensing analysis.

At the time of press, we have thoroughly tested PSF models in only the F814W filter. However, our IDL routines preserve *TinyTim*'s ability to create PSFs in other filters, and for sources of any colors. Our routines are therefore easily adaptable to other data sets.

The whole method is intentionally designed to be as adaptable as possible for many methods. The desire to know the PSF at any arbitrary position on the sky is far from unique to weak lensing. But even in lensing, advanced methods like Shapelets will, in the near future, be able to take advantage of more detailed information about the PSF shape than it is reasonable to expect from interpolation between a few stars (Massey & Refregier 2005; Refregier & Bacon 2004). This is even more exaggerated when considering higher order shape parameters, with an intrinsically lower signal to noise. The creation of noise-free, oversampled stars at any position on an image allows such analysis in any ACS data.

In the near future, we plan to release our PSF models to the community along with the wrapper we have written for *TinyTim* which will allow users to create PSF models in different filters and at user-defined positions in the ACS field of view. Interested readers are advised to contact the authors for these resources.

Acknowledgments. We would like to thank Catherine Heymans for sharing her knowledge of the ACS PSF. We are grateful to John Krist for guiding our poor lost souls through the underbelly of *TinyTim*. Great thanks go to Andy Fruchter and Marco Lombardi for useful discussions about *MultiDrizzle*. Adam Riess and Marco Sirianni provided expert knowledge about CTE effects. Richard Ellis and Alexandre Refregier have engaged us in many useful and interesting discussions about the COSMOS field. We are also pleased to acknowledge the continuing support of Nick Scoville, Patrick Shopbell and the whole COS-

MOS team in obtaining and analyzing the COSMOS images that were used to test our PSF models.

References

Bristow et al. 2002, *Modelling Charge Transfer on the STIS CCD*, in The 2002 *HST Calibration Workshop : Hubble after the Installation of the ACS and the NICMOS Cooling System*, Eds Arribas, Koekemoer, & Whitmore, STScI

Koekemoer, A. M., Fruchter, A., Hook, R. & Hack, W., 2002, 'MultiDrizzle: An Integrated Pyraf Script for Registering, Cleaning and Combining Images,' in The 2002 *HST Calibration Workshop : Hubble after the Installation of the ACS and the NICMOS Cooling System*, Eds Arribas, Koekemoer, & Whitmore, STScI, p. 337

Krist, J. & Hook, R., 2004, 'The TinyTim User's Guide' (Baltimore: STScI)

Lombardi, M. et al., 2005, ApJ, 623, 42

Massey R. & Refregier, A. 2005, MNRAS 363, 197

Massey R. et al. 2006, in preparation

Mutchler, M. & Sirianni, M., 2005, *Instrument Science Report* ACS 2005-003, (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Refregier, A., 2003, ARA&A, 41, 645

Refregier, A. & Bacon, D., 2002, MNRAS, 338, 48

Rhodes, J. et al. 2006, in preparation

Rhodes et al. 2004, ApJ, 605, 29

Rhodes, J., Refregier, A. & Groth, E., 2000, ApJ, 536, 79

Riess, A. & Mack, J., 2004, ISR ACS 2004-006,STScI

Scoville, N. et al. 2006, in preparation

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Calibration of the ACS Emission-Line Filters

C. R. O'Dell
Department of Physics and Astronomy, 1807-B Vanderbilt University, Nashville, TN 37235

Abstract. On-orbit calibration of the ACS emission-line filters F502N, F658N, and F660N have been made using a bright portion of the Orion Nebula as a reference source. This calibration is important because the conditions for pre-launch measurement of the filter profiles are not those in the ACS. The results of this work allow an absolute calibration of the surface brightness of targets in the [O III] 5007 Å, Hα 6563 Å, and [N II] 6583 Å lines. It is essential that if accurate determination of the Hα 6563 Å line is desired, one must make observations in both the F658N and F660N filters since the former can be strongly contaminated by [N II] emission.

1. Note

A complete report on this study has already been published (O'Dell 2004), and will not be repeated here.

References

O'Dell, C. R. 2004, PASP, 116, 729

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

QSO Coronagraphy and Ramp Imaging in the ACS GTO Programs

A. R. Martel[1], H. C. Ford[1], G. D. Illingworth[2], and the ACS Science Team

Abstract. We briefly discuss the processing and reduction steps of imaging data acquired in two minor modes of ACS as part of the ACS GTO programs : HRC coronagraphy of the luminous quasar 3C 273 and imaging of nearby early-type galaxies through ramp filters. We show examples of the data at different stages of the analysis. The fully reduced images reveal new morphological features, such as knots and filaments of dust and ionized gas in the ISM of the galaxies, and strong color variations in 3C 273.

1. Coronagraphy of the QSO 3C 273

3C 273 was observed with the High Resolution Channel in coronagraphic mode through the g (F475W), V (F606W), and I (F814W) filters on 19/20 Jul 2002 UT. Short target acquisition frames (10 sec through F502N) served to translate the QSO under the central 1.″8-diameter occulting spot to an accuracy of 0.5-0.8 pixels. Two orbits were dedicated to each bandpass with a CR-SPLIT of 2 per orbit. The spatial scale at the detector was binned to 0.″025 pixel^{-1}, resulting in a nominal field-of-view of $25'' \times 25''$. A reference star, HD 105281, with colors similar to those of the QSO, was observed in one orbit in the same bandpasses for accurate PSF subtraction.

The PSF of the reference star was used to remove the unresolved QSO. The PSF was manually registered and subtracted via cubic convolution interpolation until the residuals in the wings were visually minimized. The registration appears sensitive to 0.05 pixels and the largest observed shift between the PSF and QSO was 0.8 pixels. The PSF normalization was also manually adjusted until the contribution of the telescope PSF was visually minimized. Past experience indicates that this normalization method is accurate to $\approx 1\%$. The images were finally corrected for geometric distortion.

For comparison, the HRC/F606W frame before and after subtraction of the reference PSF is shown in Fig. 1. Before the subtraction (top panel), the most prominent feature is a scattered light bar, symmetric about the PSF. It is found to expand proportionately with wavelength - it might be due to large angle diffraction or a directional wavefront error. Residual diffraction rings are observed around the central spot after the PSF subtraction (bottom panel). These are caused by mismatches between the QSO and reference PSFs, specifically by differences in their positions relative to the center of the occulting spot. A color map produced from the final PSF-subtracted images is shown in Fig. 2. A detailed description of the colors and morphological features in 3C 273 can be found in Martel et al. (2003) and in the STScI press release 2003-3.

[1]Department of Physics and Astronomy, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218

[2]UCO/Lick Observatory, University of California, Santa Cruz, CA 95064

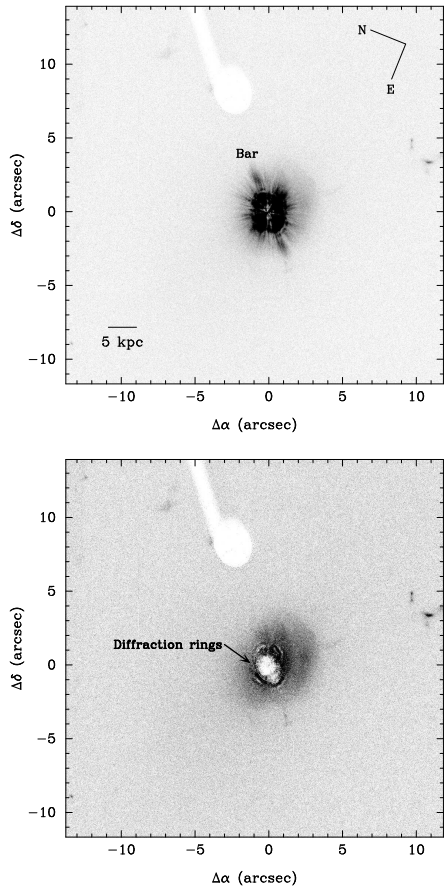

Figure 1: HRC/F606W frame of 3C 273 before (top) and after (bottom) subtraction of the QSO with a reference PSF.

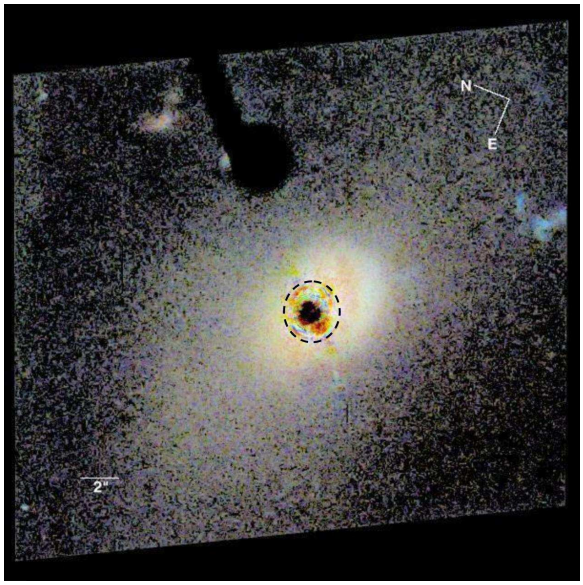

Figure 2: gVI color map of 3C 273. The dashed circle delimits the residual diffraction rings around the central occulting spot.

2. Ramp Imaging of Nearby Ellipticals

Four early-type galaxies were observed for one orbit in the FR656N filter over their redshifted Hα+[N II] emission-line complex and in a broad-band filter, F555W or F814W, adjacent in wavelength, for continuum subtraction. The broad-band exposures were limited to a total duration of 700 sec to avoid overhead penalties in the dumping of the on-board buffer memory and the remainder of the orbit was dedicated to the ramp exposure, typically 1400-1700 sec. The exposures were split into two sub-exposures for a more effective rejection of cosmic ray hits and/or hot pixels. The images were processed with the ACS Science Analysis Pipeline at Johns Hopkins University. The spatial scale of the final, reduced images is $0\farcs05$ pixel^{-1}, resulting in a nominal field-of-view of 3.5×3.5 arcmin^2.

To extract the pure emission-line maps, the intensity levels between the broad- and ramp images were matched using the counts in a small region in the host galaxy. The scaled continuum image was then subtracted from the ramp image. Ratios calculated with the SYNPHOT *calcphot* task in PyRAF with a template spectrum give very similar results. The line fluxes were measured by summing the counts in regions with positive line emission. This usually required delimiting the regions visually with polygons, boxes, and/or ellipses in SAOImage *ds9*. The conversion factor between the observed flux expressed in counts sec^{-1} and the physical flux in standard units of ergs sec^{-1} cm^{-2} was calculated with *calcphot*. The fluxes were then corrected for interstellar extinction.

A "dust map" was also generated for each galaxy by calculating the ratio of the continuum image over its isophotal model generated with the PyRAF *ellipse* task. Pixel values less than one represent dusty regions. Two examples, NGC 2832 and NGC 6166, are shown in Fig. 3. Several new, resolved dust and ionized gas features are detected in these galaxies, some on kpc scales. Further details can be found in Martel et al. (2004).

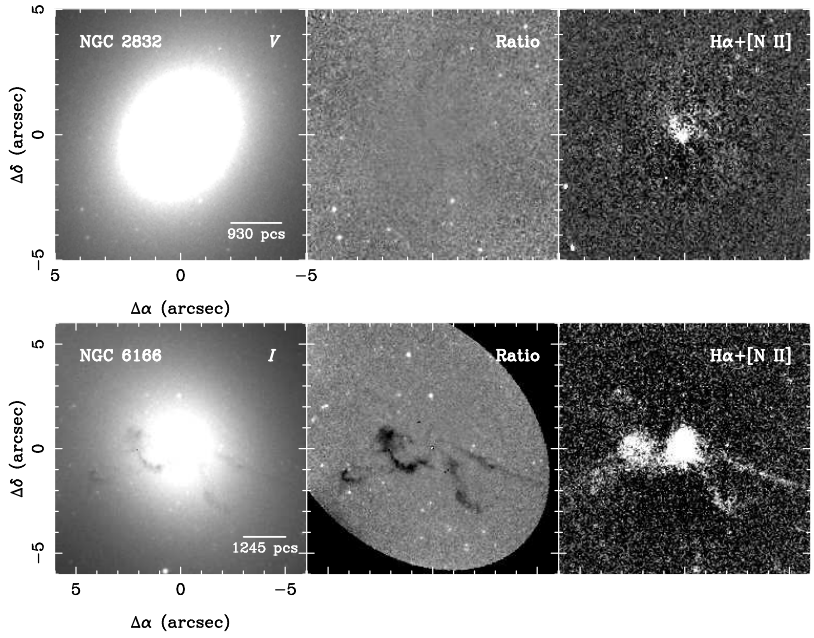

Figure 3: The continuum image, the dust map, and the Hα+[N II] emission-line map (extracted from the FR656N image) of NGC 2832 and NGC 6166 are displayed on a linear scale.

Acknowledgments. We are grateful to T. Allen, K. Anderson, and J. McCann for processing and archiving the ACS GTO data presented here.

References

Martel, A.R., Ford, H.C., Tran, H.D., Illingworth, G.D., Krist, J.E., White, R.L., Sparks, W.B., Gronwall, C., Cross, N.J.G., Hartig, G.F., Clampin, M., Ardila, D.R., Bartko, F., Benítez, N., Blakeslee, J.P., Bouwens, R.J., Broadhurst, T.J., Brown, R.A., Burrows, C.J., Cheng, E.S., Feldman, P.D., Franx, M., Golimowski, D.A., Infante, L., Kimble, R.A., Lesser, M.P., McCann, W.J., Menanteau, F., Meurer, G.R., Miley, G.K., Postman, M., Rosati, P., Sirianni, M., Tsvetanov, Z.I., & Zheng, W. 2003, AJ, 125, 2964

Martel, A.R., Ford, H.C., Bradley, L.D., Tran, H.D., Menanteau, F., Tsvetanov, Z.I., Illingworth, G.D., Hartig, G.F., & Clampin, M. 2004, AJ, 128, 2758

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

ACS Charge Transfer Efficiency. Results from Internal and External Tests

M. Chiaberge[1], A. Riess, M. Mutchler, M. Sirianni[1], and J. Mack

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. We present results of in-flight charge transfer efficiency (CTE) monitoring performed with different tests. Internal tests (Extended Pixel Edge Response [EPER] and First Pixel Response [FPR]) and external photometric test of stars in a field 6′ off the core of the globular cluster 47 Tucanae. The results of internal and external tests on parallel CTE are comparable. CTE losses increase linearly with time and they are worst at low signal levels and for low background. We provide a formula to correct for CTE losses for all flux levels, sky values and times. We also outline our plan for future monitoring of CTE.

1. Introduction

The charge transfer efficiency (CTE) per pixel is the fraction of charge transferred from one pixel to the next during readout. In an ideal CCD, the CTE would be exactly 1.0, with no charge being lost. But imperfections in the crystalline lattice of a real CCD, caused either by the manufacturing process or the space radiation environment, can act as charge traps. Although the amount of charge lost per pixel (Δ Q) is typically a very small fraction of the total charge (Q), the total CTE over N pixel transfers is CTE^N, which becomes increasingly significant as larger CCD arrays are manufactured and flown in space. Here we summarize the results of in-flight CTE monitoring performed through both internal and external tests. More details can be found in the instrument science reports (ISR) at the URL http://www.stsci.edu/hst/acs/documents/isrs.

2. Internal test

The internal monitoring programs consist of two tests performed for both the Wide Field Channel (WFC) and High Resolution Channel (HRC): Extended Pixel Edge Response (EPER) and the First Pixel Response (FPR). Since internal tests do not involve observations of real astronomical objects – the light source is always the internal Tungsten lamp (rather than stars) – they do not independently lead to a CTE correction suitable for science data. But the data can be collected with great breadth and frequency (at all signal levels, throughout the life of the detector), so relative changes are therefore useful for monitoring CTE trends.

EPER is a measurement of the excess charge found in the CCD overscans, which appears as an exponential tail following the last real pixels in the array, which tapers down to the bias level within just a few pixel transfers (see Figure 1 showing how the EPER test is performed for WFC). This tail is "deferred charge" which has been trapped during the readout, and then released on a time scale of milliseconds.

[1]Affiliated to the Space Telescope Division, European Space Agency

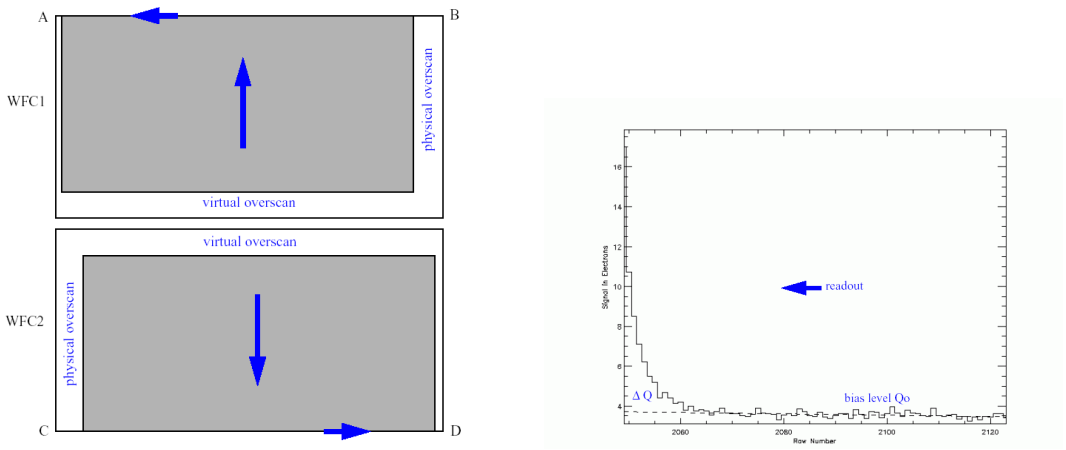

Figure 1: WFC EPER Flats with extra-large trailing overscans (75 pixels) are produced in both the parallel (vertical arrow) and serial (horizontal arrow) clocking directions. A, B, C and D indicate the location of the readout amplifiers (left panel). After subtracting the bias level Q_0, the the excess charge in the overscan is measured. This "tail" is deferred charge which has been trapped during readout, and released on a timescale of milliseconds (right panel). A similar EPER test is also performed for the HRC.

FPR measures a leading-edge loss of charge. FPR frames have a special clocking pattern where the first half of a frame is flushed (read out quickly), freshly exposing every charge trap. Then the other half of the frame is read out normally (see Figure 2 showing how FPR is used for the HRC). As the first column (or row) in the normal half of the chip is transferred across the flushed half, it loses charge as it fills most of the traps. We measure the charge lost in that first pixel, and compare it to the charge present in all the subsequent pixels in a column (or row), which suffer little or no charge loss.

2.1. Results of internal tests

- CTE loss is greatest at the lowest signal level. We observe a power-law dependence on signal level, at each epoch.
- As expected, parallel CTE is much worse than serial CTE for both WFC and HRC. No trend in time can be determined for serial CTE.
- Parallel CTE degradation for WFC is very linear at all signal level, so we can confidently extrapolate the future in-flight performance (see Figure 3). No significant difference is observed between the two chips.
- From HRC parallel EPER and FPT test we find different power laws. This may be interpreted as an "optimistic" versus a "pessimistic" estimate of the same effect, because of the different nature of the tests.

3. External test

External tests are based on observations of a field located 6′ off the center of the globular cluster 47 Tucanae. Observations have been performed twice a year during cycles 11,12 and 13. Here we present results for the WFC based on the first four epoch observations, and extrapolation to future times.

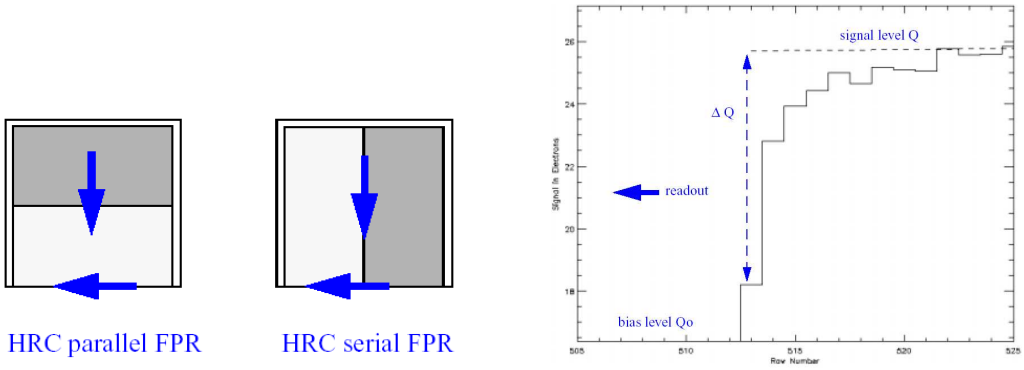

Figure 2: HRC FPR: one side of the chip is flushed (light gray) while the other side is read out normally (dark gray). Amplifier C is used for all tests (left panel). Right panel: the plot illustrates the parallel HRC test for the lowest signal level of $\sim$ 25 electrons (i.e. the worst case). A similar test is also performed for the WFC, but only for the serial register.

The goal of the observations is compare the relative photometry of individual stars as a function of the number of pixel transfers during readout.

The observational strategy is as follows: to vary the relative positions and hence the number of pixel transfers of individual stars on the WFC, we utilize 2 large scale dithers of $\sim 1/2$ the size of the WFC ($102''$) along each axis. Because of the corner placement of the amplifiers, individual stars which undergo n parallel transfer when images on chip 1, will undergo 2048-n transfers when imaged on chip 2 (see Figure 4).

Details on the data analysis can be found in Riess & Mack (2004). Here, we summarize the main points. Star list are generated from an IDL implementation of DAOfind. Stars are selected as positive fluctuations surpassing a threshold of 0.05 e^-/sec and satisfying sharpness and roundness thresholds. Regions around bright stars are masked out to avoid contamination from spurious sources. The sky background is measured as the median counts in annuli of inner radius 10 pixels and outer radius 15 pixels centered on each star.

WFC is expected to be more affected by CTE losses than HRC because of the larger number of pixel transfers. However, detailed analysis of HRC data is still in progress. For the WFC, we see strong evidence for photometric losses in the parallel direction whose value grow with decreasing stellar flux and background. No evidence for CTE losses is seen in the serial direction yet. From the analysis of stars with different count rates, different background levels and at different times Riess & Mack (2004) have derived a formula that allows users to correct photometric measurements at all times and for all combinations of stellar fluxes and background values.

$$YCTE = 10^A \times SKY^B \times FLUX^C \times Y/2048 \times (MJD - 52333)/365 \qquad (1)$$

where SKY is the background in e^-, FLUX is the flux of the star in e^-, Y is the number of transfers and MJD is the observation date (Julian day – 2400000). Coefficients A, B and C are listed in Riess & Mack (2004) for different aperture radii. For an aperture radius of 3 pixels, A=0.14, B=-0.31, C=-0.64.

Since background mitigates CTE losses, for most of the science applications which involve long integrations with efficient broad band filters the expected impact of CTE degradation is limited and manageable. The worst case arises from short integrations in narrow band filters. In a 30 sec integration with F502N, the background level is only 0.1 e^-/sec/pixel. A 20th mag star will yield few counts in this filter and the CTE losses may reach $\sim$ 0.4 mag by the end of the decade.

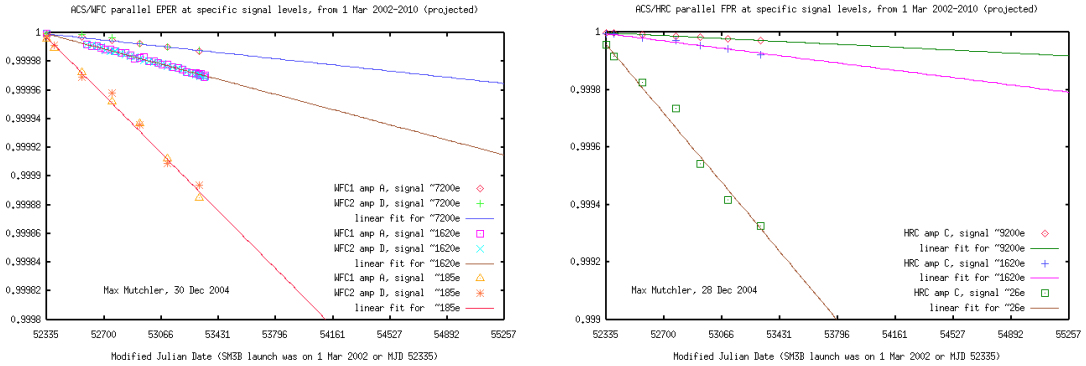

Figure 3: WFC parallel CTE from EPER test at selected signal level, over time (left panel). The linearity of CTE degradation is apparent. This gives us confidence in projecting our results into the future. Right panel: Linear decline with time at specific signal levels for HRC, from FPR test.

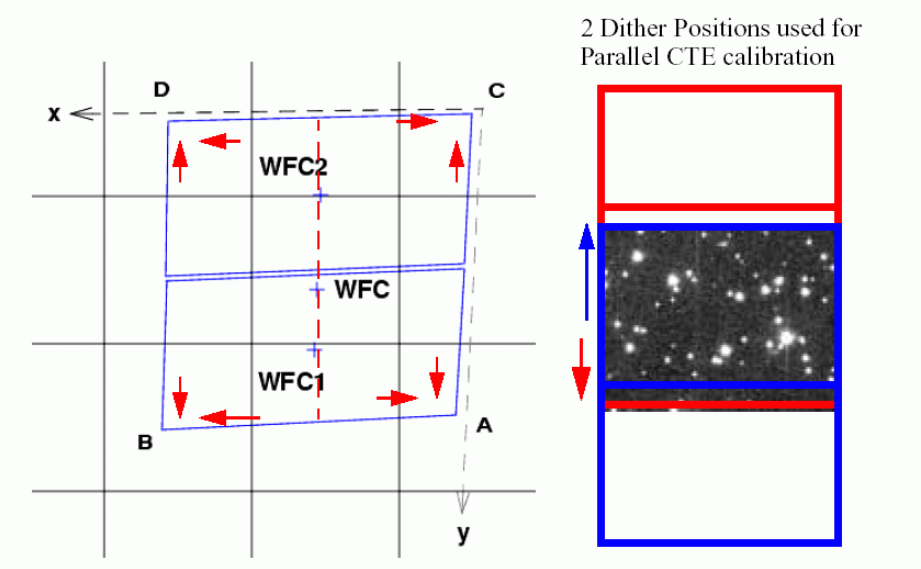

Figure 4: Schematic showing the location of the readout amplifiers and the WFC direction of readout (left panel). The right panel illustrate the two pointings used to vary the number of parallel transfers.

We conclude that for real science applications, the use of the CTE correction formula should limit the precision of flux measurements of typical faint sources to no worse than a few percent by the end of the decade.

More details on the CTE tests summarized in this paper can be found in Riess & Mack (2004), Mutchler & Sirianni (2003) and in the ACS website at the URL http://www.stsci.edu/hst/acs/performance/cte.

4. Future calibration plan

We plan to analyze the HRC data and eventually provide a formula for photometric corrections similar to that derived for WFC. New WFC data from cycle 13 and 14 will be analyzed to monitor the status of CTE losses.

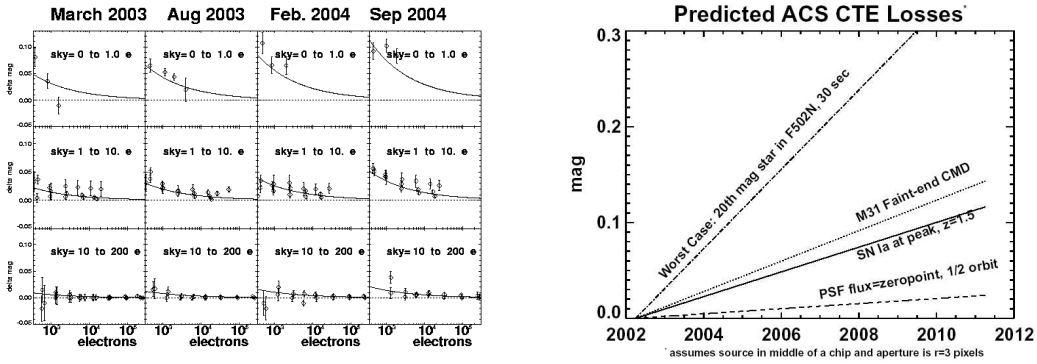

Figure 5: Left panel: photometric losses (at y=2048) for different sky levels and dates, respectively, for aperture radii of 5 pixels. The fitted line uses the global time dependent formula (1). Right panel: projected CTE losses for example science applications.

The following calibration programs aimed at monitoring CTE are being performed during Cycle 14:

- **Program 10732** Internal CTE monitoring. Data are collected at regular intervals using internal lamps only. This program emulates the ACS pre-flight ground calibration tests and post-launch SMOV tests, so that the results can be directly compared.
- **Program 10730** External CTE monitoring. Observations of 47 Tucanae with different filters and exp times. Large dithers (102″) are used to change the position of stars in the field relative to the amplifiers. This program will provide improved determinations of CTE losses for the HRC.
- **Program 10771** CTE and QE measurements for ACS instruments at different temperatures. Internal and external measurements are performed with temperatures -77, -74, -80 (WFC) and -80, -76, -85 (HRC). The external tests are performed by observing a field of the globular cluster 47 Tucanae with large dithers, as it is done in program 10730.

References

Mutchler, M. & Sirianni, M., 2005-03

Riess A. & Mack J., 2004, *Instrument Science Report* ACS 2004-006 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Determination of the CTE correction on mosaiced ACS data

E. Sabbi, M. Sirianni[1], and A. Nota[1]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. All CCDs on the Hubble Space Telescope *(HST)* suffer a progressive degradation of their charge transfer efficiency (CTE). CTE degradation can lead to photometric inaccuracy, since the measured magnitude of a star will depend on its position on the chip. Here we present a procedure to evaluate the amount of charge lost for each star in complex mosaiced Advanced Camera for Surveys (ACS) Wide Field Channel (WFC) data.

1. Introduction

All CCDs on the *HST* are subject to radiation damage, which deteriorate the detector performances (Sirianni et al. 2005). One of the most severe effects is the degradation of the CTE. Poor CTE has the effect of removing small amount of charge from the charge packets during the transfer from on pixel to the next, during the chip readout. As a consequence the accuracy of the photometry can be severely affected, since the magnitude of a star will depend, among other parameters, on its position on the chip with respect to the amplifiers. Correction to restore measured integrated counts to their "true" value is therefore required.

The charge loss due to imperfect CTE depends on the combination of many parameters, such as:

- the position of the source on the detector;
- the total counts of the source;
- the physical extent of the source;
- the level of the background;
- the amount of radiation damage suffered by the detector.

Riess & Mack (2004) provided the following CTE correction formula for WFC data:

$$YCTE = 10^A \times SKY^B \times FLUX^C \times \frac{Y}{2048} \times \frac{MJD - 52333}{365} \quad (1)$$

where coefficients A, B, and C depend on the aperture size adopted for the photometry, Y represents the number of transfers in the parallel direction, and the term $\frac{MJD-52333}{365}$ is introduced to take into account the deterioration of CTE with time. The final magnitude of a star will be: *Mag + YCTE*.

The proposed YCTE parameterization cannot be directly applied to complex mosaiced drizzled data like the one shown in Figure 1, where any given star can be at very different distances from the amplifiers, and be subject to different amount of charge losses in each of the individual frames.

[1]On assignment from the Space Telescope Division, European Space Agency

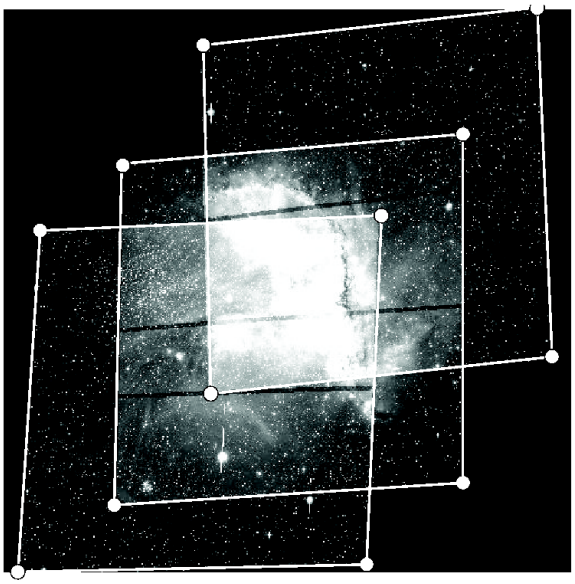

Figure 1: Mosaiced image of NGC 346. The three different pointings are shown. The positions of the amplifiers for each pointing are marked (white dots).

2. YCTE on mosaiced data

In order to derive the appropriate YCTE value for each star detected in Figure 1, we have applied the following procedure.

1. Create a deep mosaic with multidrizzle by using the long exposure frames and a more shallow one using the short exposures.

2. Star detection and PSF-fitting photometry were performed on the mosaiced images. This procedure provides the most accurate magnitude (**mag**) determination for the stars detected in the deep image.

3. Each individual frame was then corrected for geometric distortion and core aperture photometry was performed on the detected stars. This allowed us to directly measure the original integrated counts, still affected by the CTE losses.

4. YCTE correction was calculated for the photometry obtained on each single exposure, by using Equation 1 (see Figure. 2).

5. All catalogs obtained from the aperture photometry were cross-correlated and for each detected source the average value of YCTE ($<$ **YCTE** $>$) was computed for both the long and the short exposures. The single exposures are affected by cosmic rays: to obtain a clean determination of $<$ **YCTE** $>$ a sigma clipping on the associated magnitude may be required (see Figure. 3).

6. The average correction was applied to the magnitude, measured at the first step, of each star: $Mag =$ **mag**$+<$ **YCTE** $>$.

7. Given that the correction for YCTE for very faint stars is comparable to the photometric errors (see Figure. 4), stars too faint to be detected on the single exposures were not corrected for CTE.

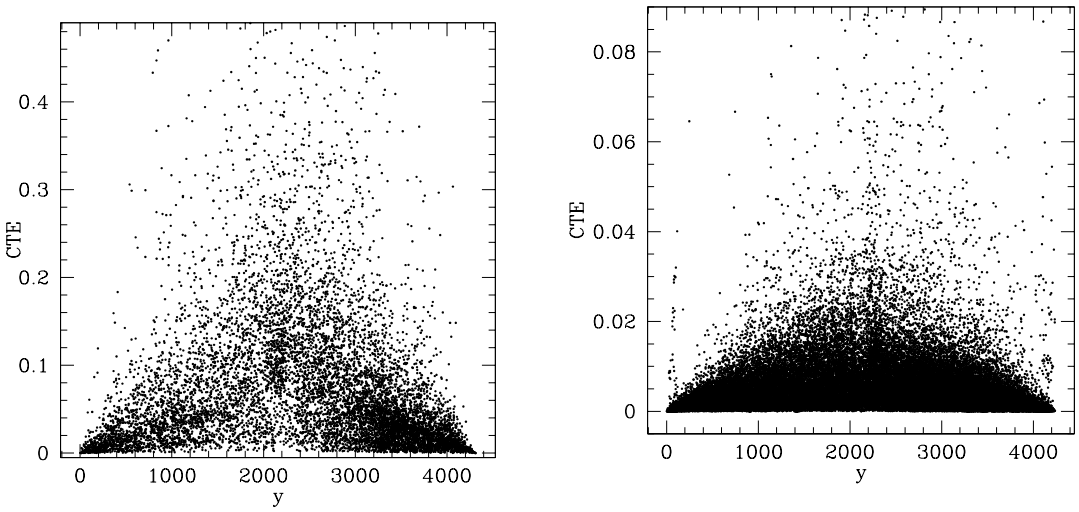

Figure 2: CTE correction as a function of the position on the frame, for a single exposure with an integration time of 3 sec (*left panel*) and with an integration time of 456 sec (*right panel*)

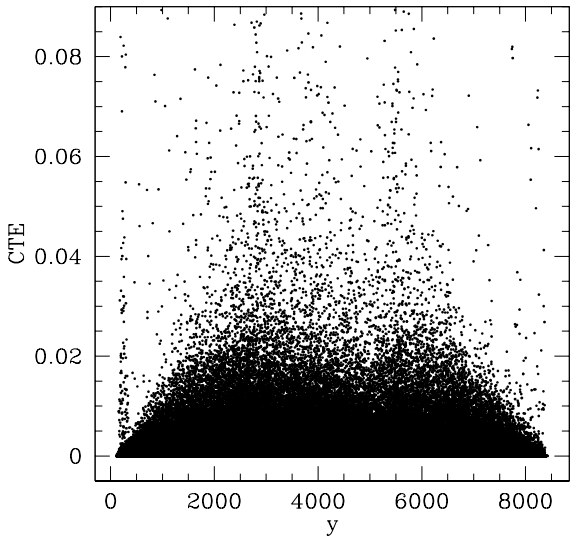

Figure 3: CTE correction for the mosaiced image shown in Figure. 1

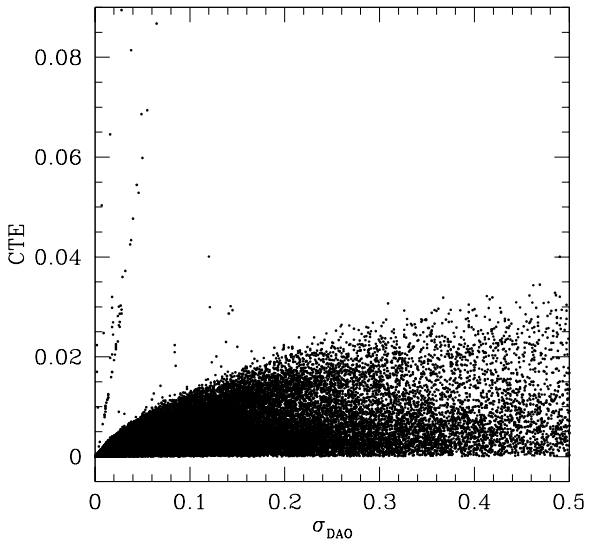

Figure 4: CTE correction for a single exposure (with an integration time of 456 sec) as a function of the photometric errors

References

Riess, A. & Mack, J. 2004, *ACS Data Handbook*, Sect. 6-17, available through http://www.stsci.edu/hst/acs

Sirianni, M., et al. 2005, PASP, 117, 1049

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Hot Pixels Growth in ACS CCDs

M. Sirianni[1]

Research and Scientific Support Department, European Space Agency

M. Mutchler & R. Lucas

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. We present an analysis of the hot pixel population in ACS CCDs. The number of hot pixels increases due to on-orbit radiation damage. In order to heal the defects responsible for the hot pixels, the CCDs are heated once a month. We review the effectiveness of the anneal process.

1. Introduction

The Hubble Space Telescope *(HST)* orbits in a Low-Earth Orbit, at about 580 km of altitude. Every day *HST* completes about 15 orbits and on average crosses some portion of the South Atlantic Anomaly (SAA) region seven times per day. The SAA is the region where the Earth's inner radiation belt makes its closest approach to the planet's surface. The population of trapped particles in the inner magnetic belt is mainly composed by protons with energy between 10 and 50 MeV, but also electrons, lower energy protons and cosmic ray ions. When *HST* transits a portion of the SAA, all of its detectors are exposed to several minutes of strong radiation. Radiation damage mechanisms in CCDs are divided in two general categories: total ionizing dose (TID) and displacement damage. Displacement damage refers to the introduction of defects in the silicon lattice. Charged particles, such as protons and neutrons can collide with the silicon atoms and displace them from their lattice sites. The vacancies created in this process can migrate in the lattice and form stable defects with another vacancy with impurities such as Phosphorus, Oxygen, and others. Any defect gives rise to a new energetic level in the bandgap and degrades CCD performance by decreasing the charge transfer efficiency, increasing the mean dark current and dark current non-uniformity by introducing individual pixels with very high dark current (also known as hot pixels and/or hot spikes).

2. Hot Pixels

Energy levels near midgap are responsible for generation of electron-hole pairs and therefore for an increase in the dark current. Some pixels show very high dark current, up to several time the mean dark rate. Depending on the particular collision sequence, protons of the same energy may produce very different amounts of displacement damage. Moreover, if a defect is created in a high electrical field region, the contaminated pixel can show very large dark current as result of field-enhanced emission. Hot pixels accumulate as a function of time on orbit. Figure 1 shows the evolution with time of the distribution of pixels at

[1]Postal Address: Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

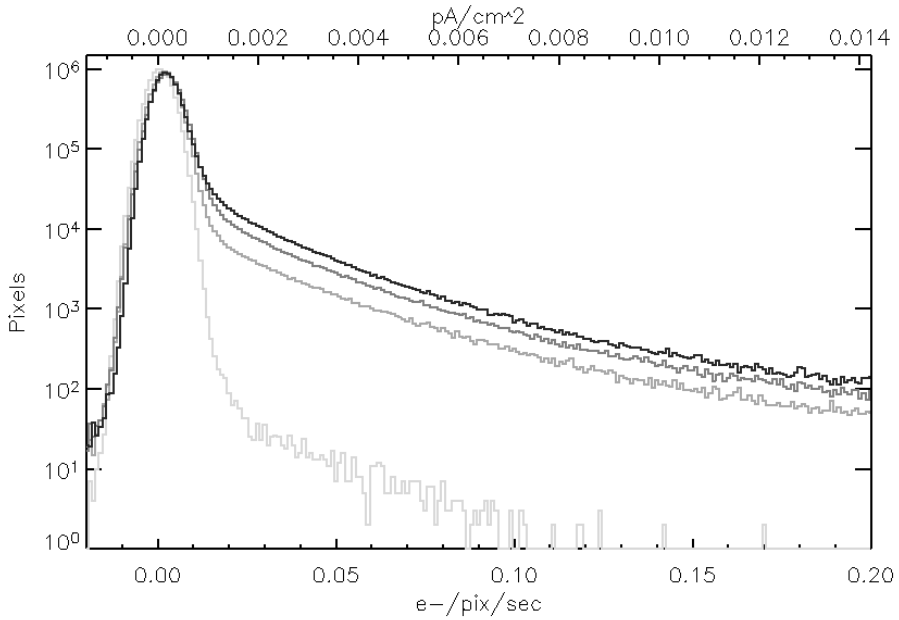

Figure 1: Distribution of dark pixels for one WFC CCD at launch (2002) (bottom), and after 1,2, and 3 years on orbit.

different dark rate for one of the ACS/WFC CCDs as a function of time. As the proton-induced damage increases, the mean dark current (the peak of the distribution) and the hot pixels population (the tail of the histogram) increase. For practical purpose we define as "hot" all pixels with a dark rate greater than 0.08 e^-/sec and as "warm" pixels whose dark current is between 5σ of the main dark distribution and the hot pixel threshold. When more than one image is available the ACS data pipeline by default rejects the pixels flagged as "hot" pixels (see Mutchler et al. 2006, this volume)

The number of pixels with a dark current higher than the mean dark current increases every day by few to several hundreds depending on their signal level (Table 1). Likely most of these new hot pixels are transient.

Table 1: Daily growth of hot pixels

e^-/pix/sec	HRC	WFC
0.02	136	824
0.04	98	607
0.08	68	489
0.10	36	390
1.00	1	16

The best approach to mitigate the impact of hot pixels in science frames is to dither the observation by one or more pixels. Only warm pixels are properly corrected after the subtraction of a "dark frame" that brackets the observation. The population of hot pixels for any given day is recorded in the Data Quality array of each image. For more on dithering strategy, visit `www.stsci.edu/hst/acs/proposing/dither`.

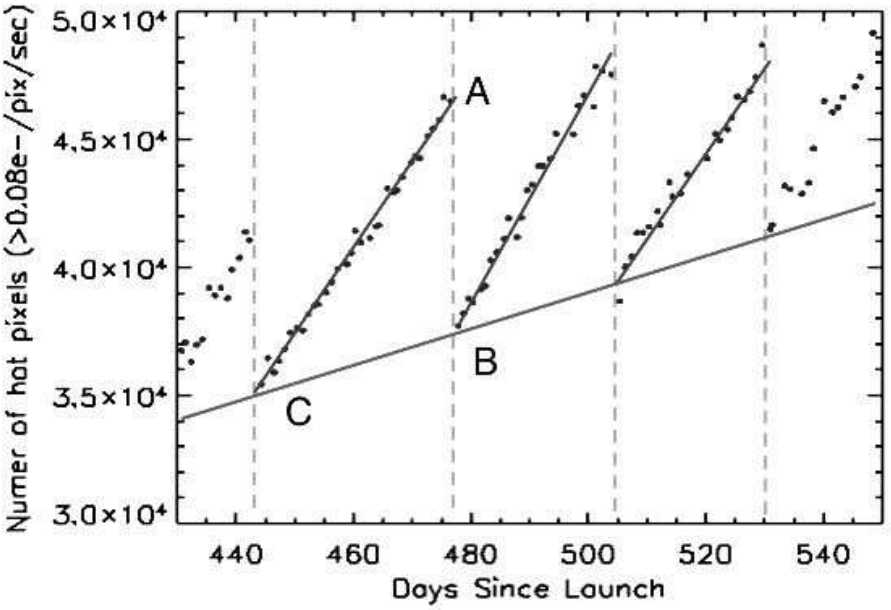

Figure 2: Determination of the anneal rate for the WFC CCDs. Dots show the population of hot pixels above the threshold of 0.08 e^-/pix/sec at each day. The vertical dashed lines mark the annealing dates. After each anneal cycle, the number of hot pixels before the anneal [A] drops to the level [B]. The ratio (A−B)/(A−C), where [C] is the population of hot pixels at the end of the previous anneal cycle, determines the anneal rate for that cycle.

3. Annealing of Hot Pixels

Like other CCDs on *HST*, the ACS detectors undergo a monthly annealing process. The TECs are turned off and the heaters are turned on to warm up the CCD to $\sim$ 19 C for a few hours. It is still not clear why at such low temperature a significant annealing is observed: the most common defects anneal at much higher temperature. By comparing the hot pixel population before and after each anneal (see Figure 2), it is possible to measure the effectiveness of the anneal process.

The annealing rate is strongly dependent on the dark current of the hot pixel (Table 2). Very hot pixels show a higher anneal rate than warm pixels and there is no impact on the average dark current level. Since launch, the duration of the CCD warming period has been reduced from 24 hrs to 12 hrs and since spring 2005 to 6 hrs. No variation in the annealing rate has been observed.

Pixels that do not anneal become permanently hot. The rate of growth of permanent hot pixels depends on the signal level. At any threshold the number of hot pixels increases linearly with time (Figure 3). By 2008 the number of pixels permanently hot will be similar to the number of pixels contaminated by cosmic rays in a 1000 sec exposure.

It is interesting to follow the status of pixels that become hot due to radiation damage. The four panels on Figure 4 show typical behaviors observed in hot pixels in ACS CCDs. In each panel, the vertical lines show the annealing dates and the horizontal line marks the threshold for hot pixel definition. The two panels on the top show two normal hot pixels become hot and resume their normal status after the first anneal cycle. The plot on the left shows a pixel that is fully healed. The one on the right is only partially healed; even after the anneal the dark current of this pixel is more elevated than normal pixels, but it remains below the hot pixel threshold. The two panels on the bottom of Figure 4 show

Table 2: Anneal Rate for ACS CCDs

	WFC		HRC	
Temp (C)	-77		-81	
Annealing Temp	-10 to +20 C		-10 to +20 C	
Annealing Time	6 to 24 hr		6 to 24 hr	
Threshold				
e-/pix/sec	%	±	%	±
> 0.02	0.55	0.02	0.64	0.02
> 0.04	0.70	0.07	0.84	0.07
> 0.06	0.78	0.04	0.84	0.04
> 0.08	0.82	0.03	0.87	0.03
> 0.10	0.84	0.02	0.85	0.02
> 1.00	0.55	0.15	0.64	0.15

pixels that, once damaged, respond to almost every anneal cycle. When they are hot the annealing can restore their status, but subsequent anneals can activate it again as hot pixel (reverse annealing).

4. Lessons learned

- The anneal rate depends on the dark current level of the pixel (Table 2).
- The number of permanently hot pixels increases linearly with time. The only solution to mitigate their presence is to dither the observations.
- The same anneal rate can be obtained at colder temperature. In at least four instances, in occurrence of *HST* safing events, the temperature of the CCDs raised to only -10 C for periods between 24 and 48 hours. After these periods the population of hot pixels decreased by the same amount as in normal cycles at +20 C (Figure 5).
- The anneal rate does not seem to be related to the length of the anneal (Figure 5).

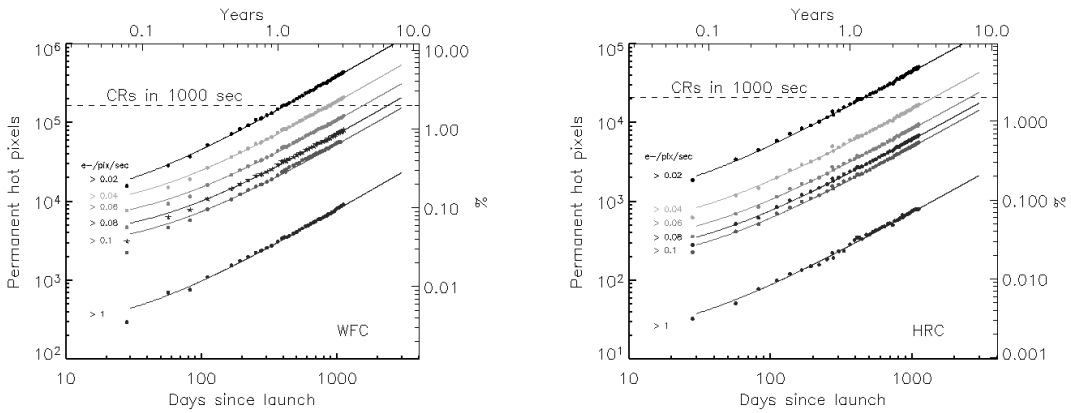

Figure 3: Growth of permanent hot pixels as a function of time for different signals.

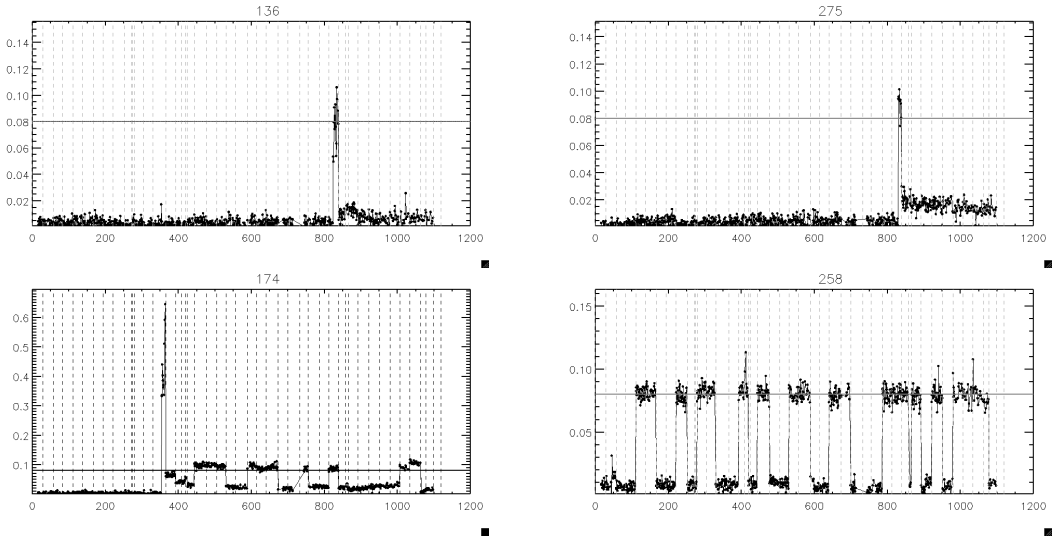

Figure 4: Dark current of selected hot pixels as a function of time (see text).

- For any particular hot pixel, a complete anneal is a rate event. Most of the annealed hot pixels significantly reduce their dark current level but never rejoin the population of normal dark pixels (Figure 4).
- Several hot pixels show evidence of reverse annealing. The process of heating the CCD or simply a power cycle can cause previously-damaged pixels to change from hot to normal and vice versa (Figure 4).

References

Mutchler, M., Sirianni, M., & Lucas, R. 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 51

Sirianni, M., & Mutchler, M. 2005, in *2005 Scientific Detector Workshop*, eds. J. Beletic & P. Amico (Garching: ESO)

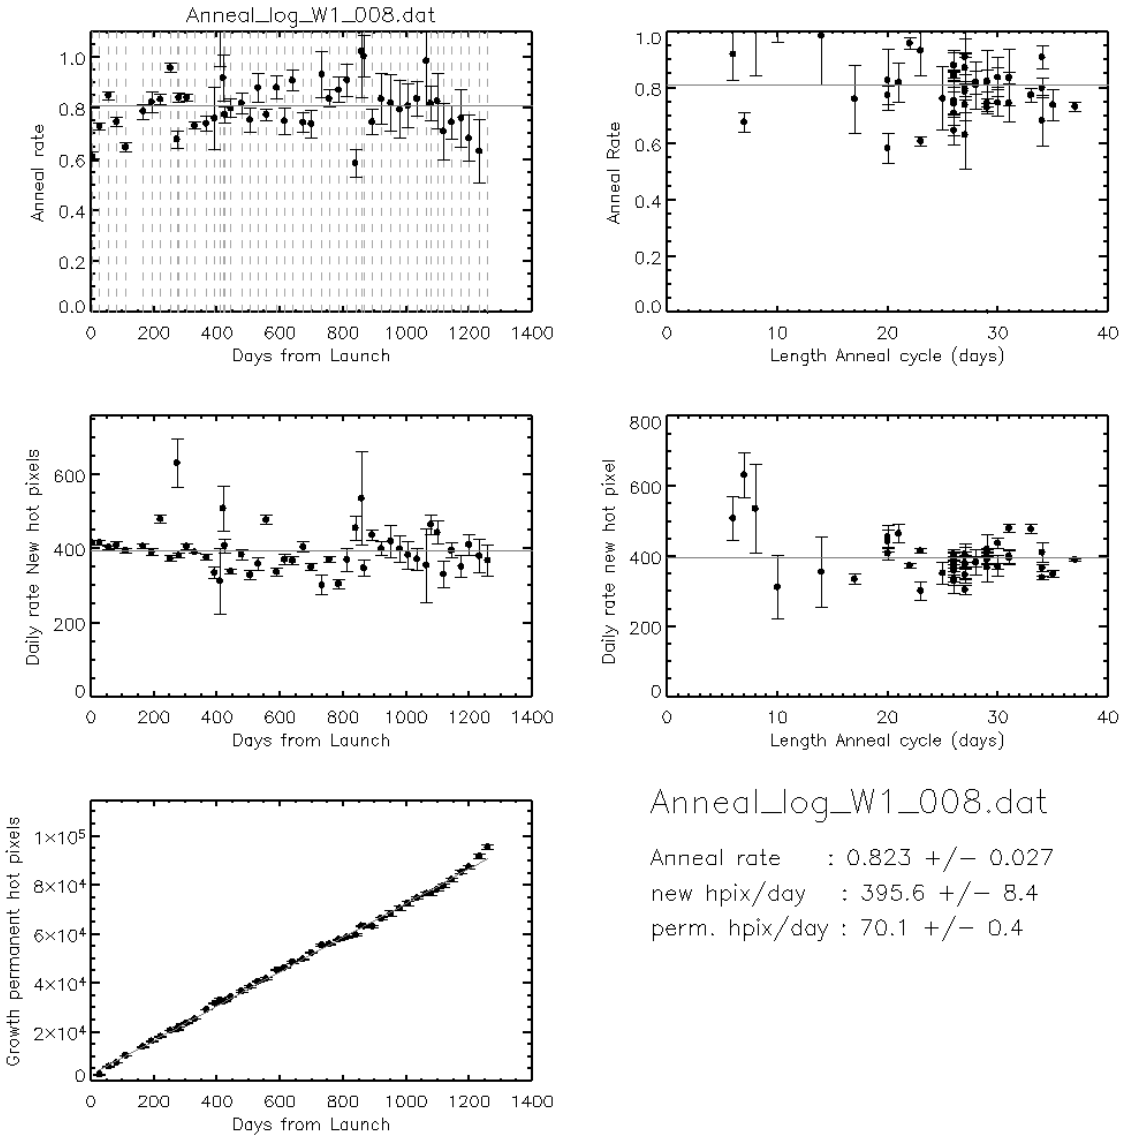

Figure 5: Example of the analysis of hot pixels for WFC. Top: Anneal rate as a function of time (left) and as a function of anneal frequency (right). MIDDLE: daily rate of hot pixel creation as a function of time (left) and as a function of the anneal frequency (right). BOTTOM: growth of permanently hot pixels.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Pipeline Calibrations of ACS Data

Max Mutchler, Marco Sirianni[1], & Ray A. Lucas

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218
E-mail: mutchler@stsci.edu, sirianni@stsci.edu, lucas@stsci.edu

Abstract. Our strategies for collecting ACS calibration data, and converting it into reference files for use in the calibration pipeline, have been continually evolving since the instrument was installed during Hubble Servicing Mission 3B in March of 2002. We provide an historical overview of basic ACS pipeline calibrations, including the impact that some of the more recent changes have on downstream data processing (e.g. drizzling). This paper emphasizes bias and dark calibrations for the ACS CCDs, and it represents a significant update to previous documentation of this subject (Mutchler et al. 2004). We describe the expected detector degradations that these calibrations are designed to track and correct, and also some unexpected anomalies we've encountered and the steps we've taken to ameliorate them.

1. Introduction: Philosophy and Practices

We have been continually improving the quality and degree of automation in producing and delivering calibration reference files for use in the ACS pipeline (**CALACS**, see Pavlovsky et al. 2005). The detector characteristics we are striving to calibrate are described in great detail in Sirianni et al. (2004). The emphasis of this paper is on bias and dark calibrations for the Wide Field Channel (WFC) and High Resolution Channel (HRC) of the Advanced Camera for Surveys (ACS). In October 2004, we implemented several improvements which we document here.

Our goal has always been to strike the best balance between quality and timeliness. We strive to produce the best calibrations possible within 2-3 weeks following any given ACS observation. For this reason, ACS users should retrieve (or re-retrieve) their data via on-the-fly-reprocessing (OTFR) several weeks after they occur, mainly to ensure that the best superbias and superdark reference files have been applied.

In August 2004, we migrated our reference file production system to the native pipeline environment (SunFire), with more streamlined production and overall quality control (Lucas et al. 2006, these Proceedings). The delivery of reference files to the Calibration Database System (CDBS) for use in the pipeline has also become more consolidated and efficient (Diaz-Miller 2005). We also deliver our reference files to the Multimission Archive at STScI (MAST) and make them available for downloading from our 'jref' directory (see Section 8).

2. Bias features and calibration

For the WFC and HRC, the bias *level* is measured from each frame's overscan region, and subtracted from each amplifier quadrant independently. Bias *features* are subtracted by the 'superbias' reference file, which is identified in the BIASFILE keyword in image headers.

[1]Affiliated to the Space Telescope Division, European Space Agency

Before October 2004, we were obtaining one bias frame every day (for both WFC and HRC), and producing a WFC superbias every week, and an HRC superbias every two weeks. Bias (mainly bad CCD columns) structure was identified (with data quality flag 128) only in a bad pixel table (BPIXTAB) produced in July 2002.

Since October 2004, we have been obtaining one bias frame every *other* day, and producing superbiases nominally every two weeks, for both WFC and HRC. The superbiases are simple cleaned combinations of the 8 bias frames obtained during each bi-week. Bias structure is now identified with flag 128 in each superbias data quality [DQ] array, which propagates to the [DQ] array of science data. Therefore, the growth of bad CCD columns is now being tracked and flagged much better, typically at two week intervals. Figure 2 illustrates the WFC bias structure.

3. Dark features and calibration

We conduct a 'monthly' (roughly every 26 days) CCD annealing which restores many hot pixels to the nominal dark current, although the population of non-annealing hot pixels grows continually (Sirianni, these Proceedings). The annealing cycle also sets the cadence for our reference file production and delivery: we make a batch of superbias and superdark reference files for every half of one anneal period (roughly every 2 weeks).

Before October 2004, we were obtaining four 1000-second dark frames every day, and producing a superdark reference file for every day. We were flagging only hot pixels in superdark data quality or [DQ] arrays, using flag 16.

Since October 2004, we began collecting 4 dark frames every *other* day, and therefore we now produce a superdark reference file for every other day. We reduced the number of bias and dark frames we obtain to lower the profile of this (the largest) ACS calibration program, with what we felt was minimal impact on the quality of the resulting calibrations. In addition to hot pixels, we also began flagging several other dark features: warm pixels, CTE tails of hot pixels, and saturated pixels, which we describe in detail in the next section.

Our dark reference files are actually a hybrid combination of two types of superdarks. First, we make a 2-week 'basedark', which is a high signal-to-noise combination of 32 dark frames. Then we also make a 4-frame 'daydark' for each day in the bi-week, which gives the best snapshot of the warm and hot pixel populations on that date. The reference superdark is a copy of the basedark, with the warm and hot pixels inserted from the relevant daydark. Figure 3 illustrates the more subtle WFC dark structures. The histogram in Figure 4 illustrate how the reference files get the majority of their pixels (normal dark current pixels, which exhibit only Poisson noise) from the 'basedark', and get their warm and hot pixels from the 'daydark'.

Due to a small bias level variation (described in Section 5), we also began equalizing any residual bias level variations seen in the superdarks, so that at least the superdarks will not propagate this problem to science data.

4. Data quality flagging and monitoring

In Table 1, we define all the flags used in the data quality arrays of ACS data (i.e. the [DQ] image extension) to identify good, bad, and corrected pixels (see Figures 5 and 6). Some of these flags emanate from the 'permanent' bad pixel tables (identified in the BPIXTAB image header keyword), although as of October 2004 we began using the bad pixel tables much less. Currently, most of these flags emanate from the generic conversion of the science data, or they propagate from the data quality arrays of the various reference files used to calibrate and combine data in the pipeline (CALACS and MultiDrizzle).

We have also re-defined some flags for new uses – warm pixels (flag 64) and the CTE tails of hot pixel (flag 32) – although in the pipeline we set parameters (bits=96) to ignore

these new flags. So for now, they are provided only as optional data processing leverage for stand-alone processing. Note that in the stand-alone environment, the MultiDrizzle default is bits=0, so the user would need to manually set bits=96 to similarly ignore these new flags (i.e., to include these pixels) during image combination.

5. Unexpected anomalies

Most of the detector calibrations were begun during pre-flight thermal vacuum testing, so their characteristics were reasonably well understood before launch. But inflight experience has revealed some unexpected anomalies. We briefly mention some notable anomalies here, and the steps we have taken to ameliorate them.

1. Random bias level variations can occur between the 'real' pixels in an image and their own overscans (Sirianni et al. 2003), where the bias level is measured, to be subtracted from the rest of the image. There is nothing we can do to correct this when it occurs in science exposures. But as of October 2004, we have been correcting this effect when seen in our superdarks (i.e. when it occurs in any of our input dark frames), so that at least our superdarks do not impose this effect on science data.
2. We have sometimes seen faint diffuse scattered light in dark frames, which would survive cosmic ray rejection and appear in the corresponding superdarks, and propagate to science data. After this problem was discovered, reference files which excluded the affected dark frames were delivered. We noticed that this problem seemed to occur following CCD annealings, when we had been leaving the filter wheels open (with CLEAR,CLEAR filters). We have not seen this problem after modifying our annealing program by rotating crossed filters into the optical path following each anneal.
3. In the first science data taken with ACS in April 2002, very faint negative imprints of bright objects were evident as mirror-image 'ghosts' replicated in each amplifier quadrant (see Figure 1). This is an electronic effect, caused by amplifier crosstalk, which shouldn't have much impact on science data. Nonetheless, we recently changed the default WFC gain (from gain=1 to gain=2) which minimizes this effect.

6. Pipeline drizzling

Although the most significant recent change to the ACS pipeline was the addition of **MultiDrizzle** for combining and cleaning associated datasets, this topic is being covered more completely elsewhere (Koekemoer 2006, these Proceedings), so we only briefly mention here some details of its implementation in the pipeline environment.

1. Pre-defined drizzling parameter sets (determined mainly by the number of images being combined) are applied to associated datasets via a FITS table, which is identified in image headers by the MDRIZTAB keyword. While this table define reasonably good default parameters for combining images in the pipeline, the user can experiment considerably with these parameters in the stand-alone environment, where trial-and-error iterations will often lead to more optimal parameters for a specific dataset.
2. As of mid-2005, pointing patterns defined with POS TARGs are now recognized as associated datasets, along with patterns defined with 'convenience' pattern forms (Mutchler & Cox 2001). Only associated datasets are automatically combined by MultiDrizzle in the pipeline. Large datasets involving multiple visits/epochs or large mosaics are not automatically or fully associated, and therefore must be combined by the user in the stand-alone environment.

3. In early 2005, filter-dependent distortion solutions (4th-order polynomials) were introduced, along with residual distortion correction images. These are identified in image headers in the IDCTAB and DGEOFILE keywords, respectively.

7. Max's wish list

The following items are some remaining issues, and thoughts on how they could be addressed. They do not necessarily represent issues for which the ACS group at STScI has given high priority. Rather, they are included here for consideration and discussion.

1. We now have two eras of data quality flagging. It would be desirable to somehow make the new flagging scheme retroactive to ACS launch (March 2002). Creating all new superbiases and superdarks for 2002-2004 and delivering them is impractical at this point. Perhaps monthly bad pixel tables could be created for 2002-2004, which reflect the new flagging scheme (contain warm pixels, CTE tails, saturated pixels). With such bad pixel tables in place in the pipeline, subsequent data retrievals (via OTFR) would have flagging much more like we have had since October 2004.
2. AS CTE worsens, it will become increasingly desirable to reject the growing CTE tails of bright artifacts such as cosmic rays and hot pixels. More complete and unique flagging of hot pixel CTE tails in superdark reference files would help with the latter. As of October 2004, the first trailing pixel is flagged uniquely (data quality flag 32), while any subsequent pixels in the CTE tail are flagged along with (i.e. not distinguished from) the warm pixels (flag 64). Note that MultiDrizzle allows the user to 'grow' the rejection of artifacts, preferentially along the trailing detector y-axis, so that their CTE tails are also rejected. However, this feature is currently not utilized in the pipeline.
3. At present, only a small fraction of ACS data is associated, so most ACS datasets are unable to take advantage of MultiDrizzle being in the pipeline. For existing archival data, logical and complete association tables could be generated offline, and ingested into the archive, such that subsequent data retrievals (via OTFR) would produce clean drizzle-combined products.
4. Add image registration (e.g. tweakshifts) to the pipeline, to allow the combination of data taken with different guide stars (e.g. data from different visits, epochs, and/or observing programs). Another way to approach this problem is to measure shifts offline, and in conjunction with the above item, ingest association tabs with the shifts and rotations embedded in them.
5. Improved handling of moving target (planetary) observations in the pipeline. Associated moving target datasets are improperly combined by MultiDrizzle, which uses the World Coordinate System (WCS) to register images. Hubble's moving target tracking is generally accurate to within a fraction of a pixel, so a combination which simply ignores the WCS would produce good results. Also, since moving targets often exhibit additional complex motions (e.g. planets rotate between exposures), single-image drizzled output is always desirable, for associated or unassociated datasets (i.e. single exposures which have only been distortion corrected).

8. More information

More information about the basic ACS pipeline calibrations described in this paper can be found on the web at: `http://www.stsci.edu/hst/acs/analysis/reference_files`.

ACS reference files in FITS format can be directly downloaded from:
`ftp://ftp.stsci.edu/cdbs/cdbs7/jref/`.

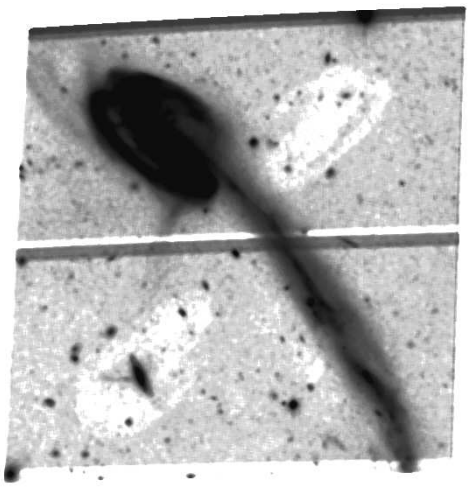

Figure 1: Amplifier crosstalk evident in the first ACS science data. The signal from the Tadpole Galaxy (in the amplifier A quadrant) generates faint negative 'ghosts' in the other amplifier quadrants.

The ACS Data Handbook and Instrument Science Reports referenced in this document are available online at: `http://www.stsci.edu/hst/acs/documents`.

Acknowledgments. We thank the following people for essential contributions to the overall success of ACS pipeline calibrations and drizzling: Warren Hack, Anton Koekemoer, Chris Hanley, Mike Swam, and Rossy Diaz-Miller.

References

Diaz-Miller, R. I., 2005, *Technical Instrument Report* CDBS-2005-01 (Baltimore: STScI), available through http://www.stsci.edu/hst/observatory/cdbs

Lucas, R. A., Swam, M., Mutchler, M., & Sirianni, M., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 61

Koekemoer A., et al., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 423

Mutchler, M., & Cox C., 2001, *Instrument Science Report* ACS 2001-07 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Mutchler, M., Sirianni, M., van Orsow, D., & Riess, A., 2004, *Instrument Science Report* ACS 2004-07, (Baltimore: STScI)

Pavlovsky, C., et al., 2005, *ACS Data Handbook*, Version 4.0 (Baltimore: STScI)

Sirianni, M., et al., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 82

Sirianni, M., Mutchler, M., Clampin, M., Ford, H., Illingworth, G., Hartig, G., van Orsow, D., & Wheeler, T., 2004, Optical and Infrared Detectors for Astronomy, Proc. SPIE, Vol. 5499

Sirianni, M., et al., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 45

Table 1: ACS data quality flag definitions

Flag	Definition
0	**Good pixels**
1	Reed-Solomon decoding error; e.g. data lost during **compression**.
2	Data replaced by fill value; e.g. neighboring cosmic ray contaminated pixels.
4	**Bad detector pixel**, or beyond aperture. In the HRC BPIXTAB, this identifies a small detector defect in the upper right corner.
8	Pixels **masked** by aperture feature, e.g. the HRC occulting finger.
16	**Hot pixels** with dark current greater than 0.08 e^-/sec; flagged in superdark data quality [DQ] arrays.
32 [a]	**CTE tails** of hot pixels; flagged in superdark [DQ] arrays. For now, we only flag the first pixel trailing each hot pixel, but this flagging may become more sophisticated and complete as CTE worsens. Note that these CTE tails are flagged more completely as warm pixels (flag 64).
64 [b]	**Warm pixels** with dark current between 0.02 and 0.08 e^-/sec; flagged in superdark [DQ] arrays.
128	Bias structure (mostly **bad columns**). Since 8 Oct 2004, bias structure has been flagged in the bi-weekly superbias [DQ] arrays. Before 8 Oct 2004, bias structure was only flagged in an increasingly outdated bad pixel table (BPIXTAB), which was created in July 2002.
256	Both full-well (useable at higher gain setting) and A-to-D (never useable) **saturated pixels** are flagged by ATODCORR, based on the CCDTAB. But A-to-D saturation is also flagged 2048, so it can be distinguished from full-well saturation. Also, since 8 Oct 2004, full-well saturated pixels in superdarks are flagged in their corresponding [DQ] arrays (note that they are also flagged as hot pixels with flag 16). Before 8 Oct 2004, saturated pixels in the superdarks were flagged only in a BPIXTAB created in July 2002.
512	Bad pixel in reference file. Used in the flatfield [DQ] arrays to indicate a portion of the flat which is not defined or not calibrated with the same accuracy as the other regions, often around the detector edges. Used for F892N and WFC polarizer observations, where the filter only subtends a portion of the chip. Used to identify dust mote replacement patches.
1024	**Charge traps**; flagged in the bad pixel table (BPIXTAB).
2048	A-to-D **saturated pixels** which are never useable, even at higher gain settings; flagged by ATODCORR, using thresholds in the CCD table (CCDTAB).
4096	**Cosmic rays and detector artifacts** rejected during MultiDrizzle image combination. These flags are not present in the drizzled output images (*drz.fits). Rather, they are propagated back to the [DQ] arrays of the input images (*flt.fits). Only data from pointing patterns (using pattern forms or POS TARGs) and CR-SPLITs are automatically associated and combined by MultiDrizzle in the pipeline.
8192	**Cosmic rays** rejected during the combination of CR-SPLIT images in the pipeline (ACSREJ); flagged in the [DQ] arrays of the output (*crj.fits) images.

[a] Re-defined as of 8 October 2004. Was previously used to flag 'blobs' of bleeding around saturated pixels (at the end of bad columns). These blobs are also flagged as hot or saturated pixels, so this was redundant.

[b] Re-defined as of 8 October 2004. Was previously used to flag 'permanent' (non-annealing) hot pixels which had persisted through the first four CCD annealing cycles after launch in 2002. These flags existed only in an increasingly outdated bad pixel table (BPIXTAB), and they were redundant with flag 16.

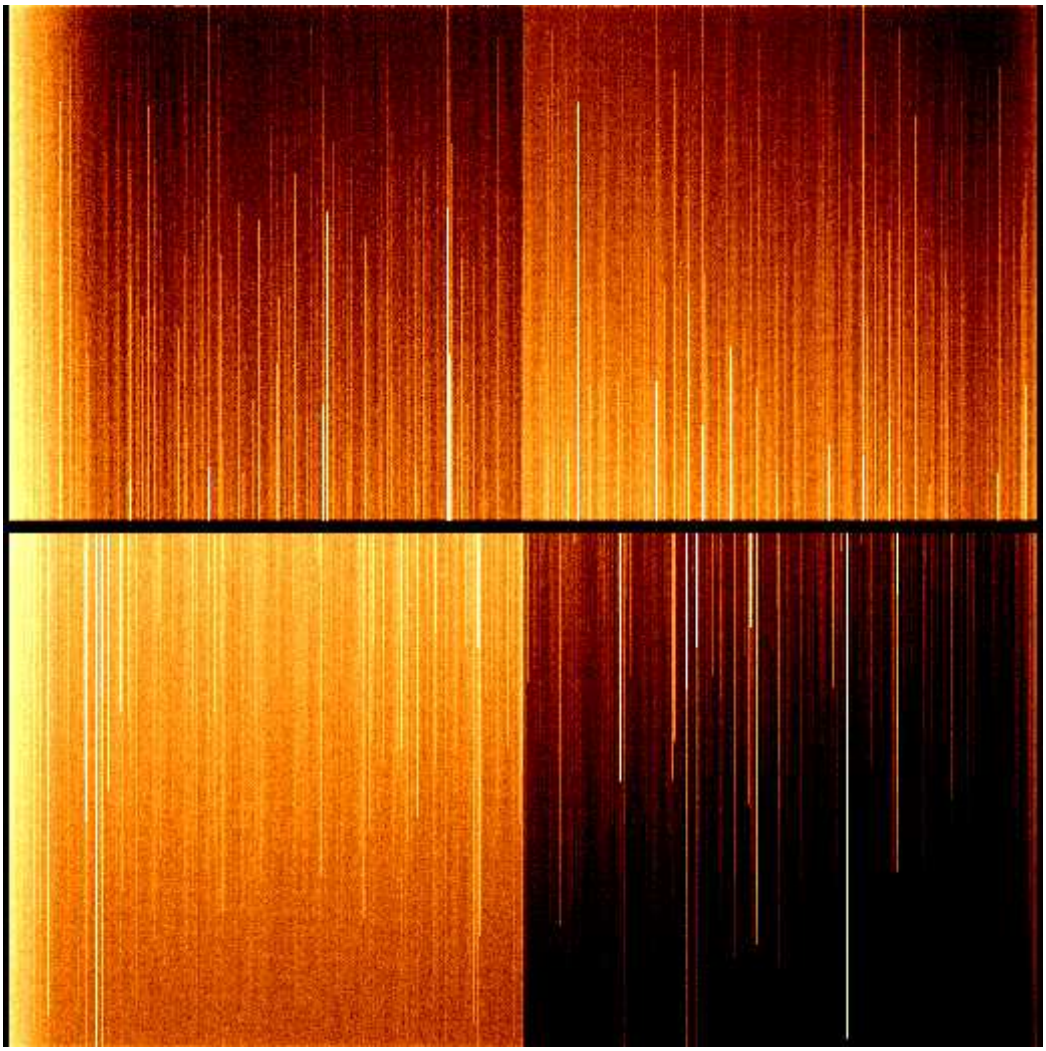

Figure 2: Bias structure in a typical WFC superbias. This image has been binned and smoothed to enhance subtle features. The two WFC chips are mosaicked: chip 1 or [sci,2] is on top, and chip 2 or [sci,1] is on bottom, with the interchip gap between them. Many bad columns are evident (their numbers are increasing with time, due to radiation damage), and each amplifier quadrant exhibits its own distinct structure. The HRC bias structure is relatively featureless, and is not displayed in this paper.

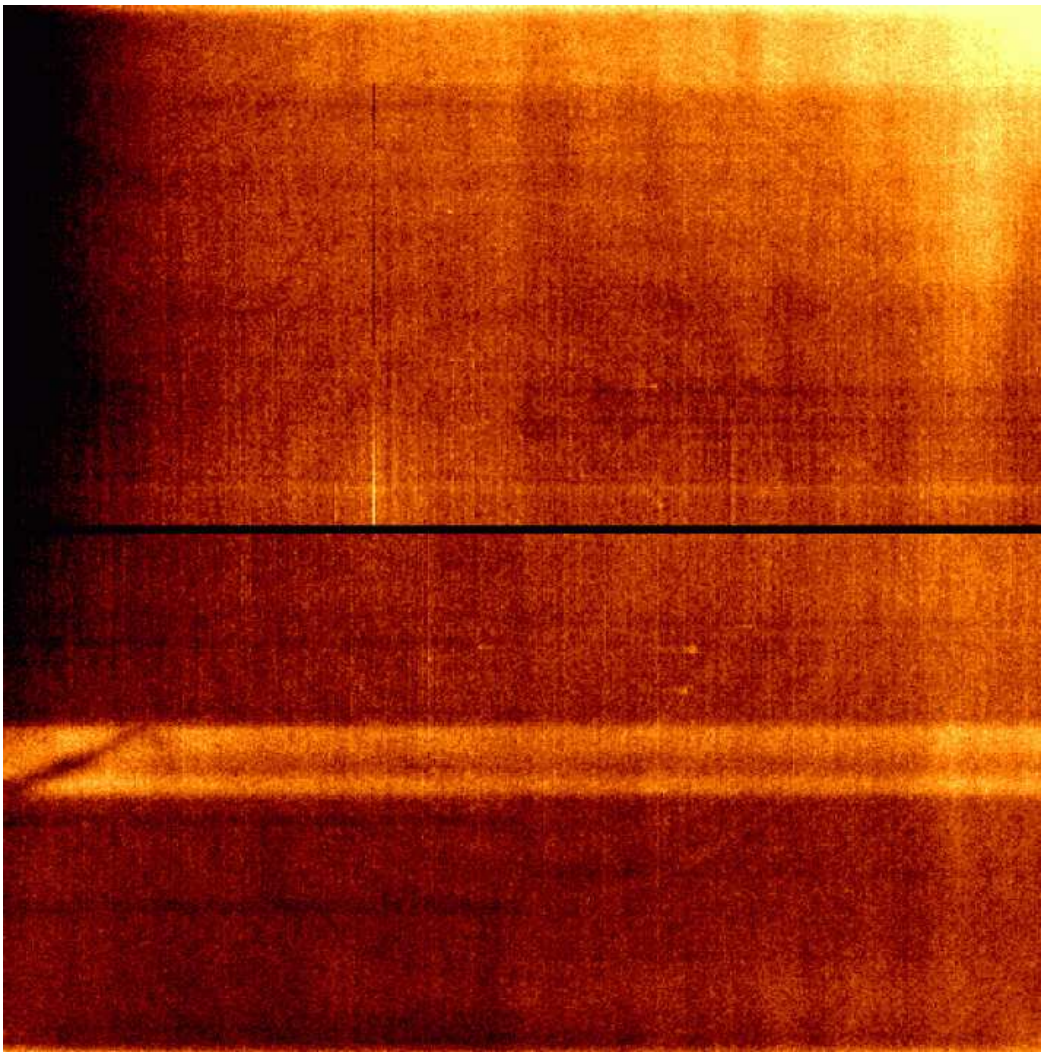

Figure 3: Dark structure in a typical WFC superdark. This image has been binned and smoothed to enhance the more subtle dark features (unobscured by the many hot pixels), and the two WFC chips are mosaicked, with the interchip gap between them. The growth of hot pixels has been well-documented (Sirianni, these Proceedings), but the more subtle structures seen here have not changed significantly since launch, and are mostly the result of CCD manufacturing processes. The HRC dark structure is relatively featureless, and is not displayed in this paper.

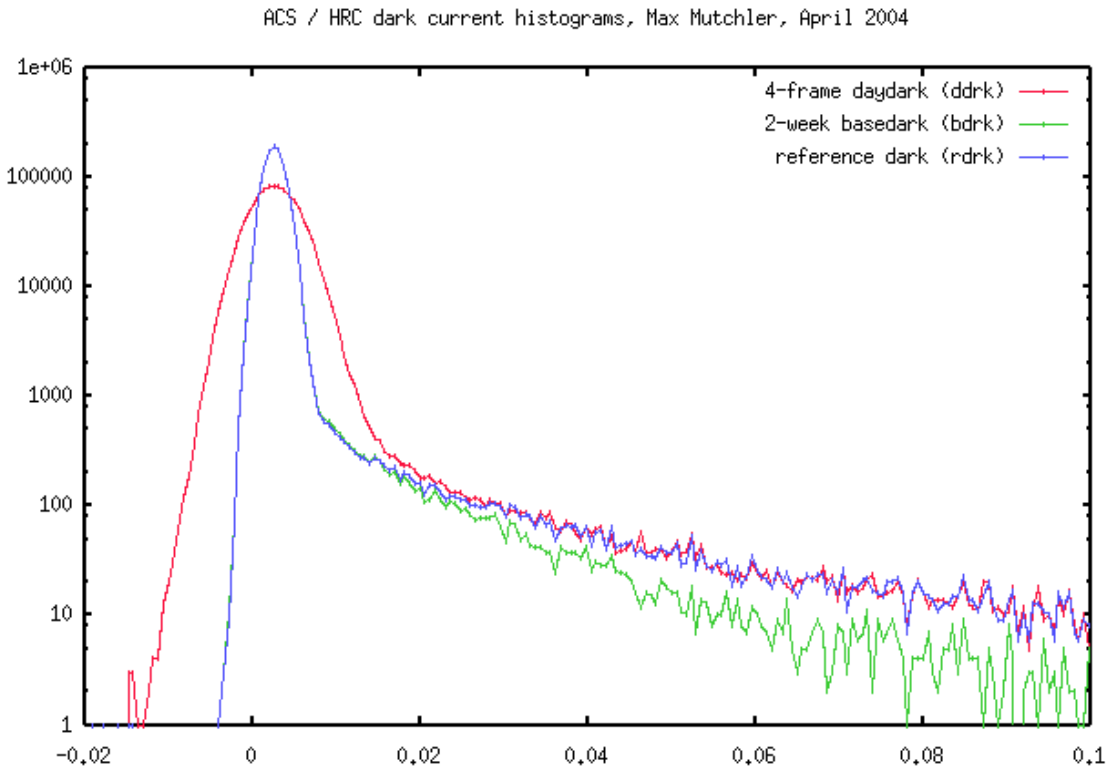

Figure 4: Overlayed HRC superdark histograms (number of pixels vs dark current in e^-/sec): a daydark, a basedark, and the resulting hybrid superdark. Our superdarks get most of their pixels from a high signal-to-noise 2-week basedark. A superdark for a given observation date is a copy of the basedark, with the warm pixels (0.2 to 0.8 e^-/sec, identified with data quality flag 64) and hot pixels (above 0.8 e^-/sec, identified with data quality flag 16) added from the corresponding 4-frame daydark.

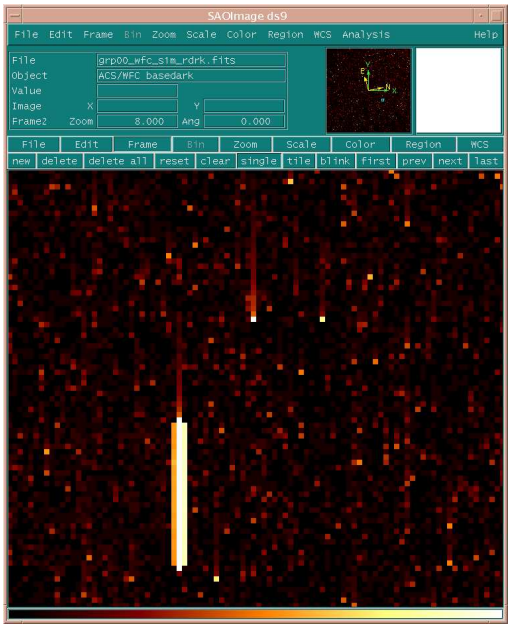

Figure 5: A closeup of some typical WFC dark structures: Scattered warm and hot pixels (some with prominent CTE tails), and columns of saturated pixels. The corresponding data quality flagging is displayed in Figure 6.

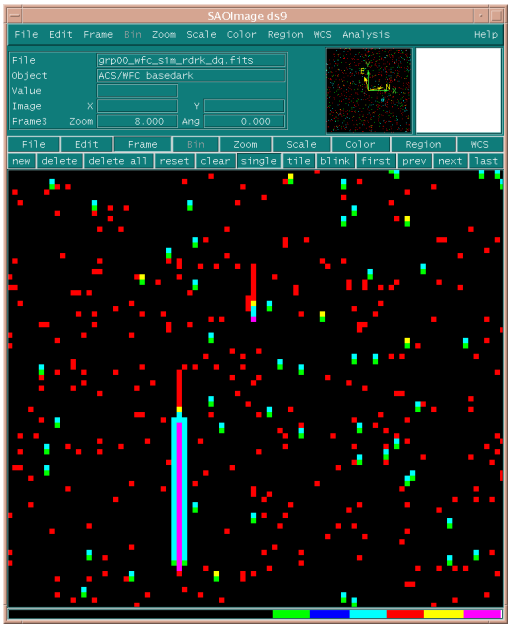

Figure 6: Data quality [DQ] flagging corresponding to Figure 5. A colorful lookup table was chosen here, to give every flag value (or combination of flag values) a unique color (which can be seen in the on-line version). Hot pixels always have one trailing pixel flagged for CTE, but more of the CTE tails are also flagged as warm pixels. Saturated columns of pixels are also flagged uniquely.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Automated ACS Bias and Dark Reference File Production, and Improved Detector Trend Analysis

Ray A. Lucas, Mike Swam, Max Mutchler, and Marco Sirianni[1]

Space Telescope Science Institute, 3700 San Martin Dr., Baltimore, MD 21218

Abstract. In the past two years, the production of ACS bias and dark calibration reference files has become almost completely automated. Building upon some code and lessons learned from the automated production of STIS reference files, the ACS bias and dark frames are automatically identified in the *HST* scheduling database and retrieved from the archive as soon as they are available on the ground. A number of formerly manual steps were also written in IRAF scripts, converted into PyRAF format, and strung together using the OPUS pipeline architecture. Routine collection of more and better image statistics has also been implemented, and these are now used in quality control along with inspection of the images. An archive of these statistics and the daily darks has been created, incorporating the files from each new biweekly period. At the same time, a series of scripts are run on the raw bias files and the daily darks to perform statistical analysis of read noise value, bias level, dark rate, and hot pixel population. The results are stored in databases that are used for longer-term trend analysis in order to better characterize the WFC and HRC cameras' performance and the effects of radiation damage, etc. as the detectors age. This document significantly updates ACS ISR 04-07 by Mutchler et al. (2004).

1. The Automated Process and Delivery of Files to OPUS and the Archive

The automated process has been derived from previous experience in building STIS CCD bias and dark reference files automatically (Swam, Goudfrooij, & Diaz-Miller 2002). It begins with the checking of *HST* observation planning and scheduling information, searching for the scheduling of ACS CCD anneal observations, and counting the number of days since the previous anneal. The scheduling and execution of an anneal triggers the formation of a biweekly period which covers all the dates going back to the boundary of the previous biweekly's data. For data taken after an anneal, the schedule-checking and counter functions determine when the next biweekly boundary (more or less halfway between anneals) will be set. If the schedule-checking monitor sees a new anneal scheduled, then the number of days since the last one is calculated and the software tries to divide the period up as evenly as possible. Alternatively, if a long enough time passes and no new anneal is seen in the schedule, the counter function counts up to a maximum number of days, creates a new biweekly boundary, and triggers the formation of a new biweekly.

At the formation of a new biweekly period, an archive request to the MAST (Multimission Archive at Space Telescope) archive is then automatically generated which requests all the raw biases and darks which were taken during the period of the biweekly. The data is automatically retrieved and copied via ftp to short-term storage and processing areas on the machine where the routine working directories reside.

[1]Affiliated to the Space Telescope Division, European Space Agency

The automated creation of the reference biases and darks then ensues, and includes these basic steps, plus the last (mostly manual) step of visually inspecting images, checking statistics, and checking for FITS compliance, etc.:

1) Combine weekly biases into biweekly superbiases for each gain, checking the rms of each amplifier quadrant, making a median-smoothed version of the superbias for visual quality assurance (QA) inspection, making a DQ (data quality) image or array of bias features with flag 128, taking statistics on bias flag counts, and updating the image header keyword content for TITLE, DESCRIP, and CCDAMP.

2) Make more binned, smoothed, mosaiced bias images and DQ mosaics for easier QA inspection, plus bias profile plots and more statistics for further review.

3) Combine two weeks of "daydarks" (multiple raw darks are taken about every other day and are combined on a per day basis to make daydarks for those days) into one higher S/N "basedark"; correct residual bias offset between quadrants; set DQ arrays to zero; set header keywords for USEAFTER and PEDIGREE, selecting dates and times from earliest raw darks to use as the USEAFTER in subsequent scripts and processing. (The USEAFTER is important for determining which is the best reference file for calibrating data based on date and time of observation, and the PEDIGREE delineates the range of observation dates of files which contributed to the superbias or superdark reference file.)

4) Make reference superdarks by adding daily hot pixels from daydarks to copies of the basedark using a preset threshold for replacement of pixels; set header keyword contents for PEDIGREE, USEAFTER, etc., and zero out dark DQ arrays again.

5) Flag warm and hot pixels, and pixels trailing hot pixels which contain deferred charge due to CTE degradation. Specifically set flags for hot pixels (16), CTE tails (32), warm pixels (64), and saturated pixels (256), adding all to DQ arrays; take statistics on dark SCI ("science" image) arrays and on DQ arrays, storing them in the long-term archive on another machine.

6) Produce more binned, smoothed, mosaiced dark images for easier inspection and still more statistics for further review. Automatically copy these additional statistics files and the daily darks to the long-term storage directory on another machine.

7) Check final image statistics, comparing to previous biweekly, visually inspect images, check header keyword content for USEAFTER, PEDIGREE, DESCRIP, TITLE, and etc., and do preliminary format checks, verifying the new superbiases and superdarks for compliance with FITS standards. If OK, give to the Calibration Database System (CDBS) group for generation of SQL and additional records for database population and delivery to the OPUS pipeline and the MAST archive for routine use and inclusion in the OPUS processing pipeline and on-the-fly-reprocessing (OTFR) requests.

A simplified schematic of the process is shown in Figure 1, and a sample OMG (process control) screen is shown in Figure 2.

2. Statistics Automatically Taken for Quality Assurance (QA)

During the routine automated processing of the reference biases and darks, we take a variety of standardized statistics which are used for QA checking of the data from the latest biweekly against that of the previous biweekly, and against that of even earlier biweeklies if necessary. These statistics, along with the daily dark images used to make the reference darks, are copied to a storage disk on another machine for use in long-term trend analysis, etc. The statistics include Npix counts for hot and warm pixels, CTE tails, and the like for both SCI arrays or images (biases and darks) and their DQ arrays. We also visually inspect the images.

Figure 3 shows a tkdiff screen with side-by-side comparison of statistics from two different biweeklies. Figure 4 shows a typical binned, smoothed, mosaiced (to enhance subtle features for further quality assurance checking) bias and dark image, respectively.

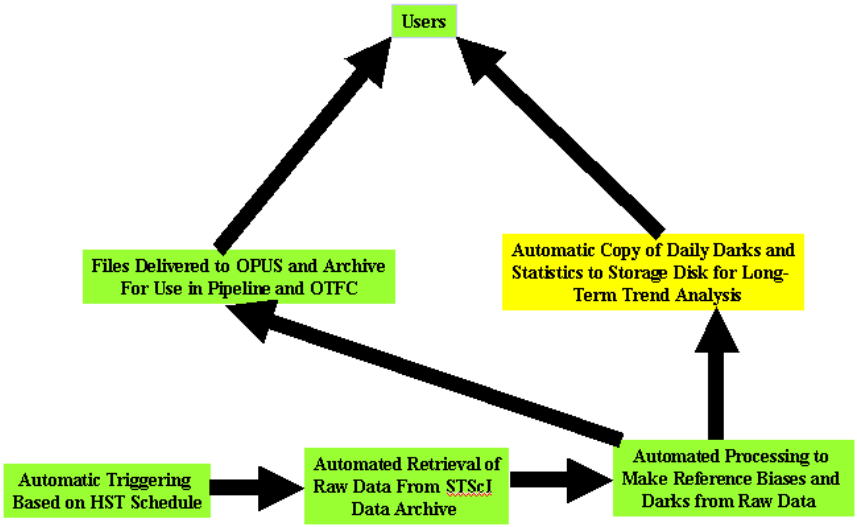

Figure 1: Schematic showing how the automated process supports users.

3. Long-Term Monitoring and Trend Analysis

As an integral part of processing, the daily darks and a variety of statistics collected on the data are copied to a storage disk on another machine. There, a suite of IDL scripts are run on the raw bias files and daily darks to measure bias level, read noise, dark current, hot pixels, and cosmic ray contamination. Two different databases are populated - the first for bias level and read noise, and the second one for the dark current and hot pixels. Graphical routines read the databases and allow a quick comparison of the statistics for the current biweekly dataset and all previous data. Hot pixel growth can therefore be well-monitored as well as the overall trend in dark current. The comparison with historical datasets not only allows us to have an instantaneous snapshot of the detectors' condition, but also allows us to identify various changes in a timely manner. For example, in June 2004, the read noise in WFC amplifier A increased by almost 1 e^- (rms). Timely discovery of similar problems is essential to plan new calibrations and to notify users of significant problems.

Figure 5 shows the aforementioned example of read noise jump in WFC Amp A, and the concurrent read noise for WFC Amp B which shows no read noise jump. Figure 6 shows the total number of WFC1 hot pixels versus day and anneal dates (vertical dashed lines) since launch, and Figure 7 shows the dark rate evolution versus time for the HRC CCD.

References

Mutchler, M., Sirianni, M., Van Orsow, D., & Riess, A., 2004, *Instrument Science Report* ACS 04-07, http://www.stsci.edu/hst/acs/documents/isrs/isr0407.pdf

Mutchler, M., Sirianni, M., & Lucas, R. A., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 51

Sirianni, M., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 45

Swam, M., Goudfrooij, P., & Diaz-Miller, R. I., 2002, in Astronomical Data Analysis and Software Systems XI, ASP Conf. Proc. Vol. 281, eds. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: Astronomical Society of the Pacific), 277

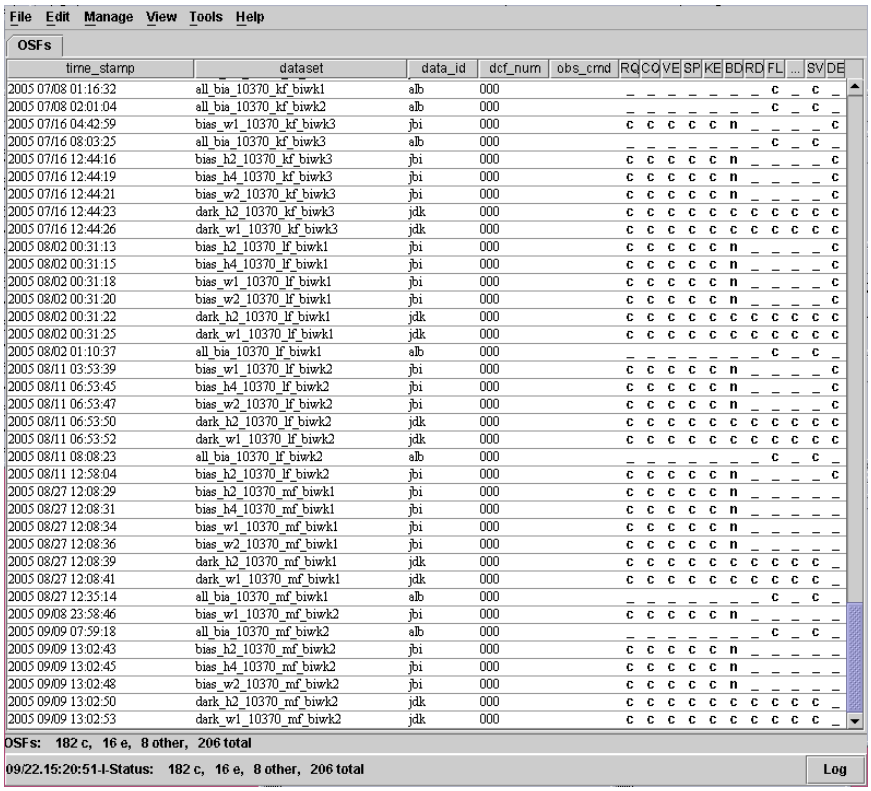

Figure 2: A sample OMG (process control) screen.

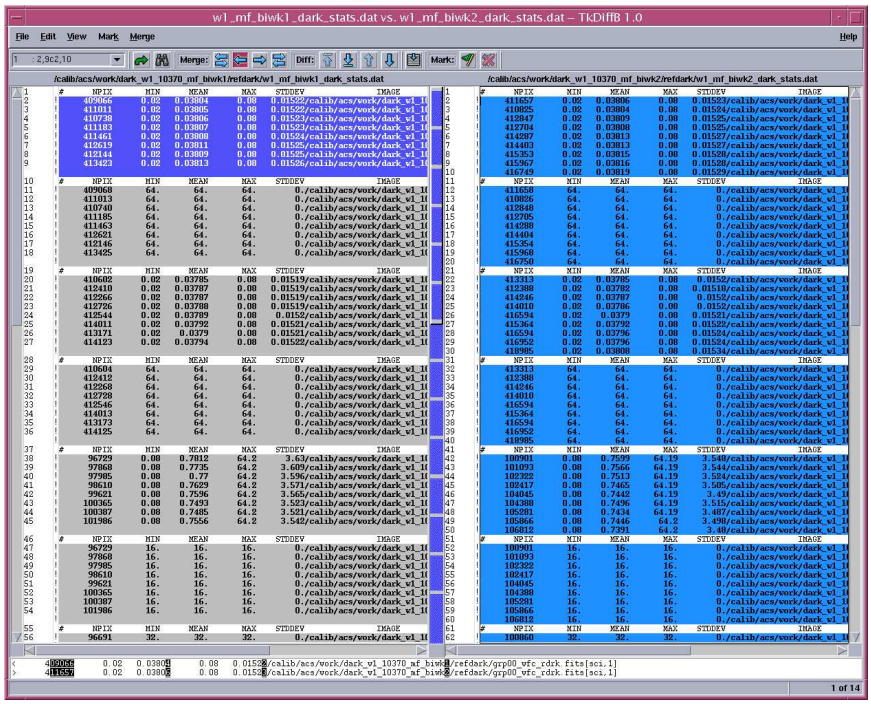

Figure 3: A tkdiff screen showing a side-by-side comparison of statistics for successive biweeklies.

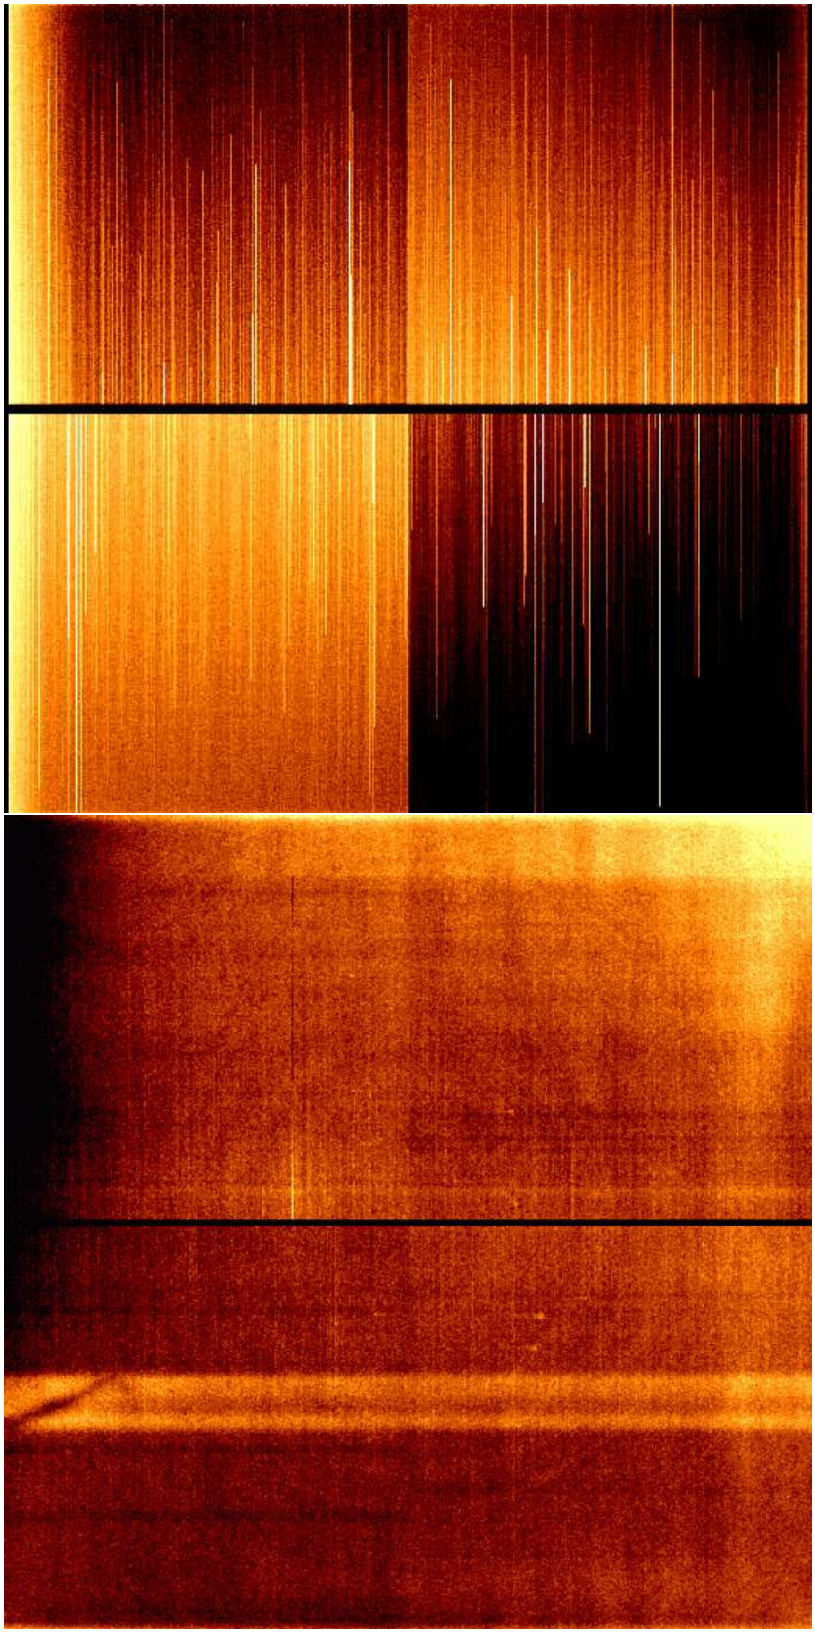

Figure 4: Superbias (top) and superdark (bottom) for the ACS Wide Field Channel (WFC). These images are binned and smoothed to enhance subtle features. The two CCD chips are also mosaiced to show their relative orientation, including the interchip gap and overscans.

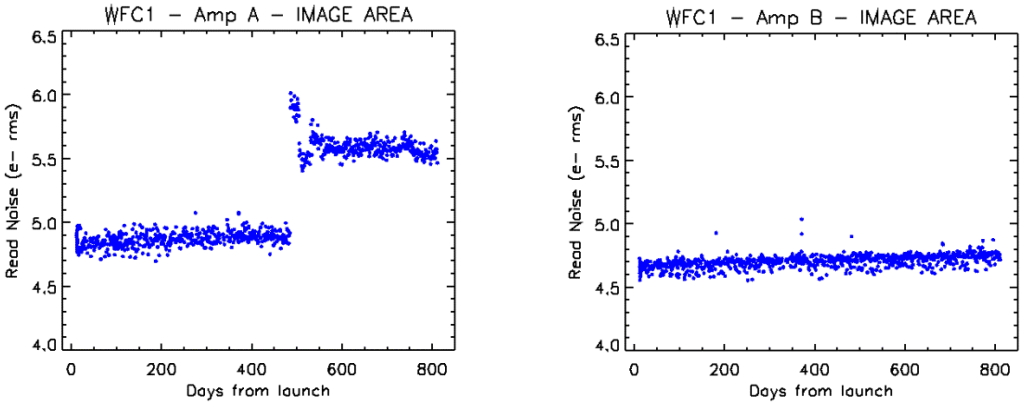

Figure 5: An example of read noise jump in WFC Amp A and no concurrent read noise jump in WFC Amp B.

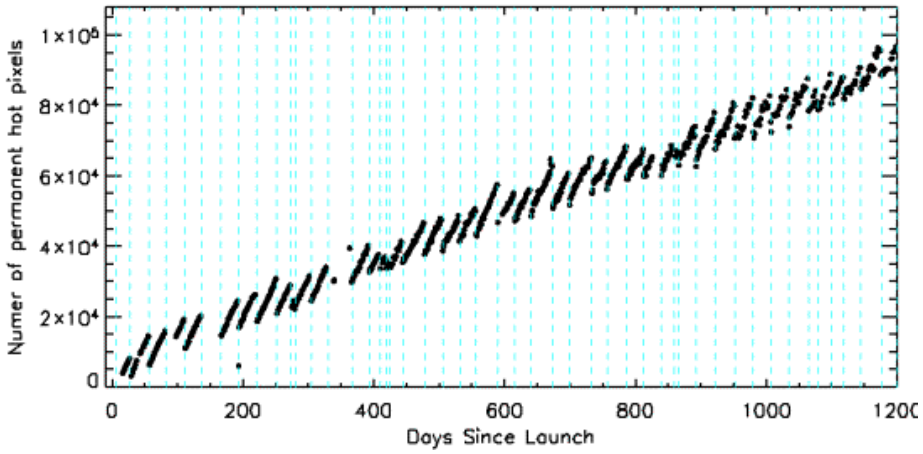

Figure 6: WFC1 total hot pixels per day since launch, showing anneal dates with vertical dashed lines.

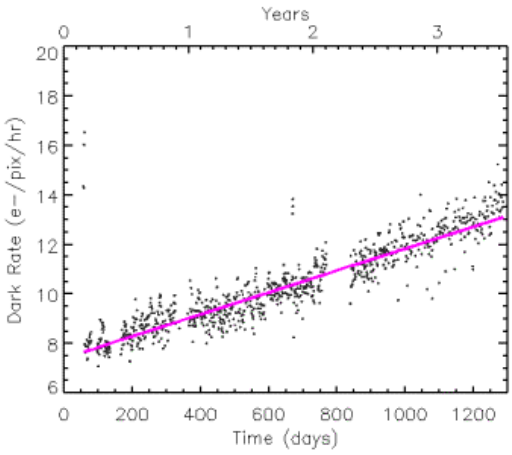

Figure 7: Evolution of dark rate over time for the HRC CCD.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

ACS Flat Field Update and New SBC L-flats

J. Mack, R. C. Bohlin, R. L. Gilliland, and R. P. van der Marel

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. We present an overview of the methods used to create the flat fields which are currently in use by the ACS calibration pipeline. L-flats for the SBC detector have recently been computed using dithered star cluster observations. Large corrections to the existing flats are required at the 10-20 percent level, depending on wavelength. A comparison with sky flats obtained via observations of the bright Earth indicate that the current pipeline flats are accurate at the 1-2 percent level.

1. Introduction

The flat field reference files currently in use by the ACS calibration pipeline were derived via different methods, depending on the detector and the filter used. These LP-flats are a combination of two files, a P-flat which accounts for the pixel-to-pixel variations in sensitivity, and an L-flat which models the low-frequency variations in sensitivity over the detector field of view. The CCD P-flats were created in the laboratory using an external illumination source, while the P-flats for the SBC were taken inflight using the internal lamps. The L-flats were derived from dithered stellar observations of 47 Tuc for the CCDs and NGC 6681 for the SBC MAMA. The two files were then multiplied to make the LP-flats used for calibration.

This paper consists of several topics. Section 2 describes the techniques used to derive the LP-flats and their uncertainties for each detector/filter combination. Section 3 discusses the new SBC L-flats, which have been recently been computed and delivered to the calibration pipeline. Section 4 describes using Earth flats as an independent means of validating the stellar L-flat technique and gives an estimate of the expected errors in the existing flats. Finally, Section 5 describes known artifacts which may be present in calibrated images and which may be misinterpreted as flat fielding defects.

2. Flat fielding Overview

At visible wavelengths (greater than ∼4000 Å), the CCD pipeline flats were created from a combination of P-flats taken in the laboratory prior to launch (Bohlin et al. 2001) and L-flats derived inflight from dithered observations of 47 Tuc (Mack et al. 2002; van der Marel 2003). From the stellar photometry, local variations in the detector response of ±3% to ±9% were found, where the required correction to the lab flats increases with wavelength. These corrections were derived from observations in 6 broad-band WFC filters and 8 broad-band HRC filters, and are labeled '47 Tuc' in Table 1. The accuracy of the LP-flats derived with this technique is expected to be better than 1%.

Assuming a simple linear dependence on wavelength, the L-flat correction for the remaining broad, medium, and narrow-band filters was computed. Using the pivot wavelength of each filter for the interpolation, these L-flats were determined from the weighted average L-flat of the two filters nearest in wavelength. These filters are labeled 'Interp' in Table 1 and are expected to be accurate at the $\sim 2-3\%$ level. More than 3 years of followup

observations of the same field have been obtained and will be used to improve the accuracy of the interpolated flats and to quantify any temporal changes in the L-flats while in orbit. These data will also provide a means of investigating any loss in absolute sensitivity as the detectors age.

For the HRC UV filters, 47 Tuc does not provide enough signal to adequately determine the L-flats, so observations of the bright Earth were used to create the pipeline LP-flats (Bohlin & Mack 2003). While the Earth is a poor flat field source at optical wavelengths due to structure in the cloud cover, the Earth is a uniform source of diffuse light below ∼4000 Å due to the high optical depth above the cloud layer. Unfortunately, red leaks in F220W and F250W are so large that the out-of-band light dominates, and the laboratory flats are superior to the observed Earth flats. Hence, large variations in the low-frequency response may still exist for these filters. Dithered observations of NGC 6681 have recently been obtained with the HRC UV-filters will be used to quantify the L-flat corrections required for these two filters. The HRC F330W and F344N pipeline flats, on the other hand, are defined entirely by the Earth flat data. With ∼20-30 observations each, over the course of 3 years, these flats have very high signal-to-noise and show repeatability to much better than the required 1% accuracy.

Table 1: CCD L-flat Summary

Filter	WFC Pivot Wave	L-flat Method	WFC Total Correction	HRC Pivot Wave	L-flat Method	HRC Total Correction
F220W	n/a	n/a		2255	None	
F250W	n/a	n/a		2716	None	
F330W	n/a	n/a		3363	Earth	
F344N	n/a	n/a		3434	Earth	
F435W	4317	47 Tuc	16%	4311	47 Tuc	7%
F475W	4744	Interp		4776	47 Tuc	7%
F502N	5023	Interp		5023	Interp	
F555W	5360	47 Tuc	10%	5356	47 Tuc	6%
F550M	5581	Interp		5580	Interp	
F606W	5918	47 Tuc	14%	5888	47 Tuc	7%
F625W	6311	Interp		6296	47 Tuc	7%
F658N	6584	Interp		6584	Interp	
F660N	6599	Interp		6599	Interp	
F775W	7693	47 Tuc	13%	7665	47 Tuc	8%
F814W	8060	47 Tuc	15%	8115	47 Tuc	9%
F892N	8915	Interp		8916	Interp	
F850LP	9055	47 Tuc	18%	9145	47 Tuc	12%

The SBC MAMA flat fields were derived inflight, using internal observations of the deuterium lamp (Bohlin & Mack 2005). This illumination does not simulate the OTA optics and, therefore, does not accurately model the low frequency variations in sensitivity. Thus, the internal flats are useful only for correcting the pixel-to-pixel detector response. As done for the CCDs, the SBC L-flats were derived using dithered starfield observations. Instead of 47 Tuc, the UV-bright globular cluster NGC 6681, which is rich in blue horizontal branch stars, was selected. While the data were taken just after launch, they were not analyzed until recently, due to the infrequent usage of the SBC detector. Due to the failure of STIS in August 2004, the SBC is now the only HST instrument with FUV capabilities, and the ACS team has shifted the SBC calibrations to high priority.

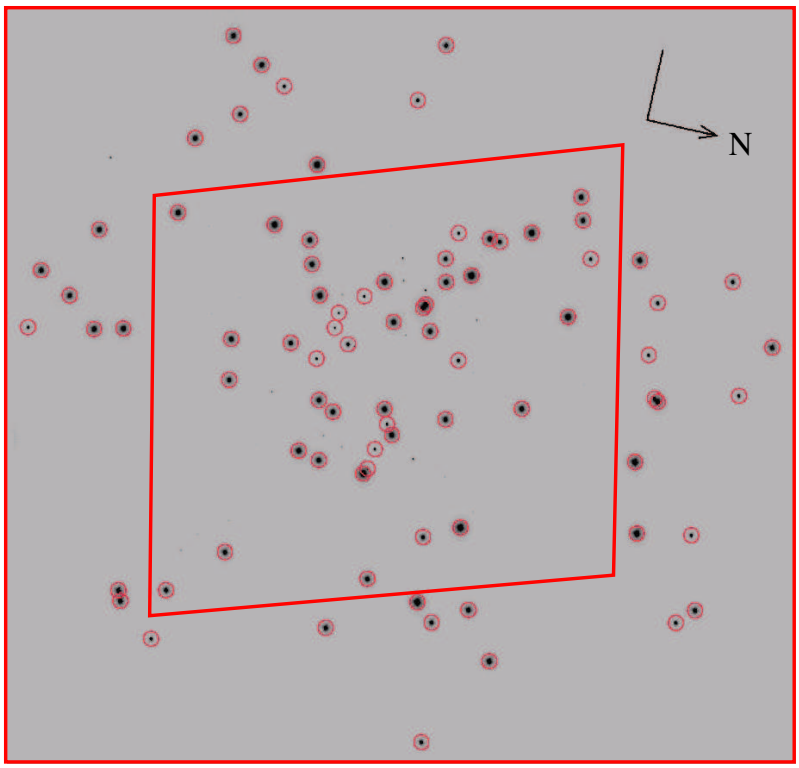

Figure 1: NGC 6681 mosaic from the combined, drizzled SBC exposures. Note that this mosaic is significantly larger than the normal SBC FOV (shown in red at the center of the figure) due to the combination of dithered observations. The 'rhombus-shaped' field of view is a result of correcting for the SBC geometric distortion. Stars used to derive the improved flat fields are circled in red. The compass at top-right indicates the direction of north, where the length of the vector is 5″. The coordinates at the field center are RA=18:43:12.75, Dec=−32:17:32.76.

3. SBC L-flats

The uniformity of the SBC detector response has been assessed using multiple dithered observations of the globular cluster NGC 6681 (Mack et al. 2005). The drizzled mosaic of this field is shown in Figure 1. By placing the same stars over different portions of the detector and measuring relative changes in brightness, low-frequency variations in the detector response (L-flats) have been derived and used to create improved SBC flat fields for the FUV imaging filters. Follow-up observations of the same field have been used to investigate any temporal dependence in the detector sensitivity.

The matrix-solution program developed for computing the CCD L-flats (van der Marel 2003) has been adapted to handle the SBC data. To summarize, the observed magnitude of a star at a given position is assumed to be the sum of the true magnitude plus a correction term, the L-flat, that depends on the position on the detector. When a set of dithered observations is available for a given star field, the determination of both the L-flat and the true instrumental magnitude of each star can be written as an over determined matrix equation. This equation has a unique minimum χ^2-solution that can be efficiently obtained through singular-value-decomposition techniques.

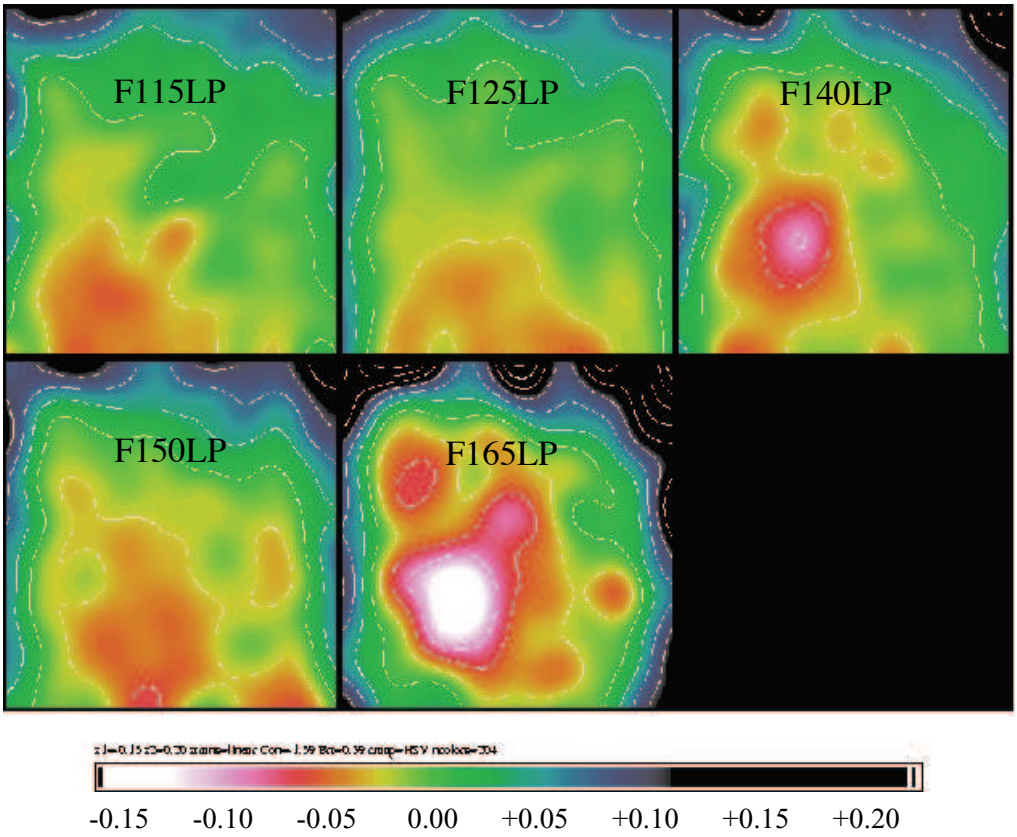

Figure 2: SBC L-flat corrections, in magnitudes, where the contours correspond to 0.03 steps in magnitude. Positive magnitudes (black) indicate that the photometry obtained using the existing P-flats was too faint, and negative magnitudes (white) indicate that the photometry was too bright.

The required low-frequency corrections to the inflight lamp flats are summarized in Table 2, and the smoothed L-flat images are shown in Figure 2, where the corrections range from ±8 to ±20 percent, depending on wavelength. Insufficient data were available to derive a solution for F122M, so the L-flat is simply a copy of the F115LP L-flat which is closest in wavelength. Errors of several percent may still exist for the F122M flat field.

In addition to the low-frequency spatial variations in sensitivity, a time-dependent component has also been found, suggesting a decline in the absolute UV sensitivity of ∼2 to 4 percent per year for the first 1.6 years in orbit. This is consistent with the situation for the STIS FUV-MAMA, where a loss of 1 to 3 percent per year was measured for the first 5 years in orbit (Stys, Bohlin, & Goudfrooij 2004). Since that time, the SBC sensitivity appears to have stabilized. Additional followup observations are ongoing.

The new SBC flat fields have been delivered for use in the calibration pipeline, and the resulting photometric accuracy is now ±1% for F115LP, F125LP, and F140LP and ±2% for F150LP and F165LP, after correcting for both the new L-flats and the time-dependent sensitivity. Further development to ACS pipeline software is required before the new sensitivity corrections can be applied within the *HST* on-the-fly-reprocessing (OTFR) pipeline.

Table 2: SBC L-flat Summary

Filter	SBC Pivot Wave	L-flat Method	SBC Total Correction
F115LP	1406	NGC 6681	16%
F125LP	1438	NGC 6681	17%
F140LP	1527	NGC 6681	23%
F150LP	1611	NGC 6681	21%
F165LP	1758	NGC 6681	40%

4. HRC Earth Flats

In order to verify the accuracy of the pipeline flat fields, several hundred observations of the bright Earth during occultation were obtained using the full set of HRC standard filters (Bohlin et al. 2005). While most of these images show streaks or other non-uniform illumination, a significant number are defect-free and can be used for comparison with the existing pipeline flats. Unfortunately, the WFC Earth flats suffer from a shutter light leak, and the Earth limb is too bright for SBC observations, so the HRC is the best detector for comparison.

In Table 3, the total number of streak-free flats is given in column 2 and the total range of deviation (minimum to maximum) between the pipeline LP-flat and the new average Earth flat is given in column 3. In general, the flat fields are confirmed to a precision of $\sim 1\%$, validating the stellar L-flat technique. While the 'interpolated' L-flats were expected to have larger uncertainties, they do not appear to be significantly worse than the L-flats derived from direct observations. One exception is the F550M filter which shows a total deviation of more than 2%. Other exceptions are the four longest wavelength HRC filters which show large systematic differences with the pipeline flats. These differences appear to be caused by stray light originating from the detector surface, where most of the long wavelength photons are reflected and then scattered back from nearby focal plane structures. Any filter transmitting at these long wavelengths will see the extra pattern from this light, though the strength of the additional stray light will be proportional to the total flux of the source. Thus, for large diffuse objects that fully illuminate the detector, these Earth flats are more appropriate for calibration than the existing pipeline flats, which are appropriate for point sources.

Table 3: HRC Earthflat Observations

Filter	Nobs	Total Deviation	L-flat Method
F330W	20	0.4±0.7%	Earth
F344N	34	0.1±0.9%	Earth
F435W	10	0.9±0.9%	47 Tuc
F475W	6	0.7±0.8%	47 Tuc
F502N	6	1.3±0.4%	Interp
F555W	4	1.8±1.6%	47 Tuc
F550M	4	2.4±1.3%	Interp
F606W	3	1.1±1.1%	47 Tuc
F625W	6	1.4±0.8%	47 Tuc
F658N	6	1.2±0.7%	Interp
F660N	3	0.9±1.0%	Interp
F775W	2	2.6±0.5%	47 Tuc
F814W	3	4.0±1.9%	47 Tuc
F892N	6	6.5±1.5%	Interp
F850LP	4	8.4±3.0%	47 Tuc

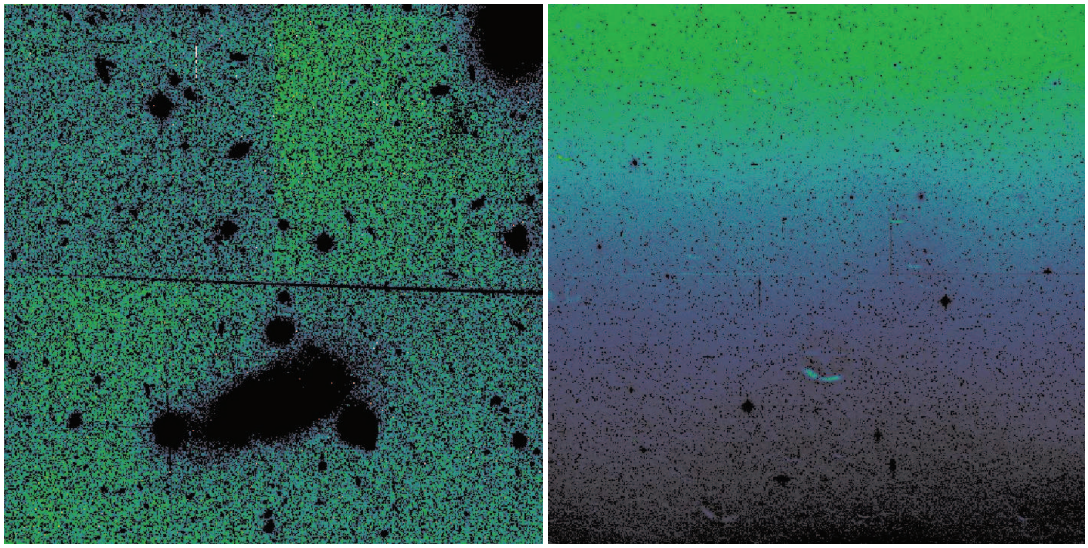

Figure 3: WFC bias offset (left) and scattered Earth light (right).

5. Artifacts affecting the 'flatness' of calibrated observations

5.1. Bias Offsets

Calibrated WFC science images processed with CALACS often exhibit a residual offset in their absolute levels at the boundary of the A-B and C-D amplifiers. In Figure 3 (left), a calibrated WFC F814W image, binned 8x8, shows this quadrant-dependent residual bias offset between amplifiers A and B. This effect can be attributed to uncertainties in the bias level subtraction due to random variations in the difference between the leading physical overscan and the bias level in the active area (Sirianni et al. 2003), where the amplitude of this residual offset is typically less than 4 DN. The total background level in each science image is therefore composed of the actual sky background, modulated by the spatially-varying quantum efficiency of the detector, plus a bias offset that is constant over the image quadrant/amplifier. Failure to remove this additional offset prior to flat fielding will result in calibrated images which appear to have a slight flat fielding error. This effect will be more apparent in narrow-band images which have a lower contribution from the sky background.

5.2. Scattered Earth Light

In an effort to minimize the background light in HST images, observations are nominally scheduled only when the telescope is more than 20 degrees from the bright Earth limb. The total background contributed by scattered Earth light depends on the angle between the HST pointing and the nearest Earth limb, and whether the limb was sunlit (bright) or dark. This additional background light is not constant, but appears as a gradient in brightness across the field of view. This gradient may be misinterpreted as poor flat fielding, while in fact it is a true representation of the sky background. In Figure 3 (right), an extreme example of scattered light is shown for an Earth limb angle of 14 degrees. This WFC F606W image is binned 8x8 and shows a gradient of 15% across the detector. Such scattered light is best removed by fitting the sky background with a two-dimensional surface, rather than assuming a constant sky value.

References

Bohlin, R. C., Hartig, G., & Martel, A. 2001, *Instrument Science Report* ACS 2001-11 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Bohlin, R. C. & Mack, J. 2003, *Instrument Science Report* ACS 2003-02 (Baltimore: STScI)

Bohlin, R. C. & Mack, J. 2005, *Instrument Science Report* ACS 2005-04 (Baltimore: STScI)

Bohlin, R. C., Mack, J., Hartig, G., & Sirianni, M. 2005, *Instrument Science Report* ACS 2005-12 (Baltimore: STScI)

Mack, J., Bohlin, R. C., Gilliland, R. L., van der Marel, R. P., Blakeslee, J. P., & de Marchi, G. 2002, *Instrument Science Report* ACS 2002-08 (Baltimore: STScI)

Mack, J., Gilliland, R. L., van der Marel, R. P., & Bohlin, R. C. 2005, *Instrument Science Report* ACS 2005-13 (Baltimore: STScI)

Sirianni, M., Martel, A. R., Jee, M.J., van Orsow, D. & Sparks, W. B. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 82

Stys, D. J., Bohlin, R. C., & Goudfrooij, P., 2004, *Instrument Science Report* STIS 2004-04 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

van der Marel, R. P. 2003, *Instrument Science Report* ACS 2003-10 (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The ACS Solar Blind Channel calibration

Colin Cox

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. The ACS MAMA detector, also known as the Solar Blind Channel (SBC) is currently HST's principal detector in the far UV. Since the loss of STIS, more attention is now being paid to SBC calibration. New dark images have been supplied and the flat fields formerly measured on the ground have been refined. The prism modes have also been recently calibrated and an improved distortion correction is being prepared.

1. Description

The ACS Solar Blind Channel (SBC) is a spare STIS MAMA operating in the far-UV from 1150 to 1700 Å. It contains a CsI photocathode deposited directly on a curved microchannel plate. Electrons generated by photons striking the photocathode enter the microchannel plate which sends an amplified electron cloud onto a multi-anode plate configured so that the incident position can be decoded. The SBC has high resolution (pixel size of 0.03 arcsec) and sensitivity plus extremely low dark current ($\lesssim 10^{-5}$ counts per pixel). The performance and principal calibrations were reported at the HST Calibration Workshop in 2002 (Tran et al.). Here we describe what changes have been made or are planned.

2. Flat Field

The flat-field image illustrates various features of the SBC. The general ribbed pattern of the anodes and the hexagonal bunching of the micro-channels can be seen. Although disconcerting in this strongly stretched image, the effects of these features are eliminated by the pipeline flat-field processing. Slightly more troublesome are the non-functioning rows due to a broken anode and the partial obscuration of several columns by the repeller wire. Because of the missing rows, we have recently (November 2004) moved the target reference point from the central row 512 to row 400 so as to reduce the likelihood of a target landing on the unresponsive region. The SBC flats, originally developed from ground-based measurements, have been recently revisited and refined incorporating in-flight data. A report on this topic is given at this workshop (Mack 2006) and there is a recent related ACS Instrument Science Report (Bohlin & Mack 2005)

3. Dark Rate

The SBC dark rate, although low, is temperature dependent. The SBC temperature is not controlled. The temperature and hence the dark rate depend principally on how long the high voltage has been switched on. A series of dark measurements led to the graph in Figure 3. The data points indicated by solid markers are part of a contiguous series of measurements. It is seen that the dark rate starts to increase fairly rapidly when the temperature exceeds about 24°C. In normal operation this only happens when the high voltage has been turned on for about three hours. The original dark current image included

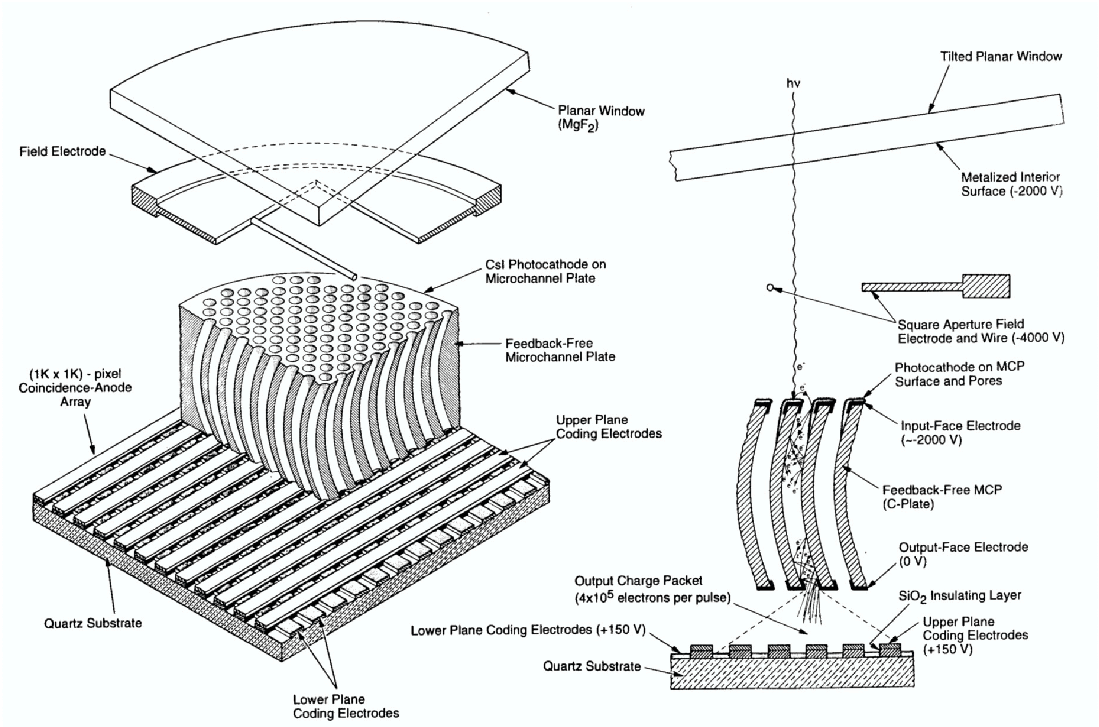

Figure 1: Image production in the ACS SBC.

all the data illustrated. The calibration file has been replaced by one which uses only data from observations taken when the detector temperature was not above 22°C. This results in a mean dark rate of 8.0×10^{-6} counts per second per pixel, whereas the previously supplied file had a count rate of 4.0×10^{-5}. The raw data show typically one or two counts per pixel with many pixels having zero counts. Using these data directly would imply in many cases that one pixel has a dark current while its neighbor does not. To provide a statistically more meaningful dark image the data have been smoothed. Although the dark rate shows clear temperature dependence, we have not embarked on a program to supply a series of darks at different temperatures. The low rate, while desirable for standard observations, makes it very time-consuming to collect enough counts for a useful dark image. The SBC is normally turned off and observations which keep the SBC on for less than three hours will be well served by the supplied dark image. The dark rate will, in any case, normally be small compared with the background and statistical noise.

4. Distortion

The SBC distortion has not so far received the same attention as have distortions in the WFC and HRC. The current polynomial fit is based on that of the HRC. The rationale for this is that the HRC and SBC are in almost identical optical paths, the only difference being a flat fold mirror which selects the active channel. The SBC solution is derived directly from the HRC with a modification of the scale, a displacement of a few arcseconds in the field of view and a small rotation. These relative adjustments were derived from back-to-back observations of the same field in each detector. Small distortions in the detector itself have been measured and are on file, but have not been applied to the published calibration. A distortion analysis of the STIS MAMAs is under way and will be extended to include the ACS MAMA. The detector distortions, having very localized features, could not be

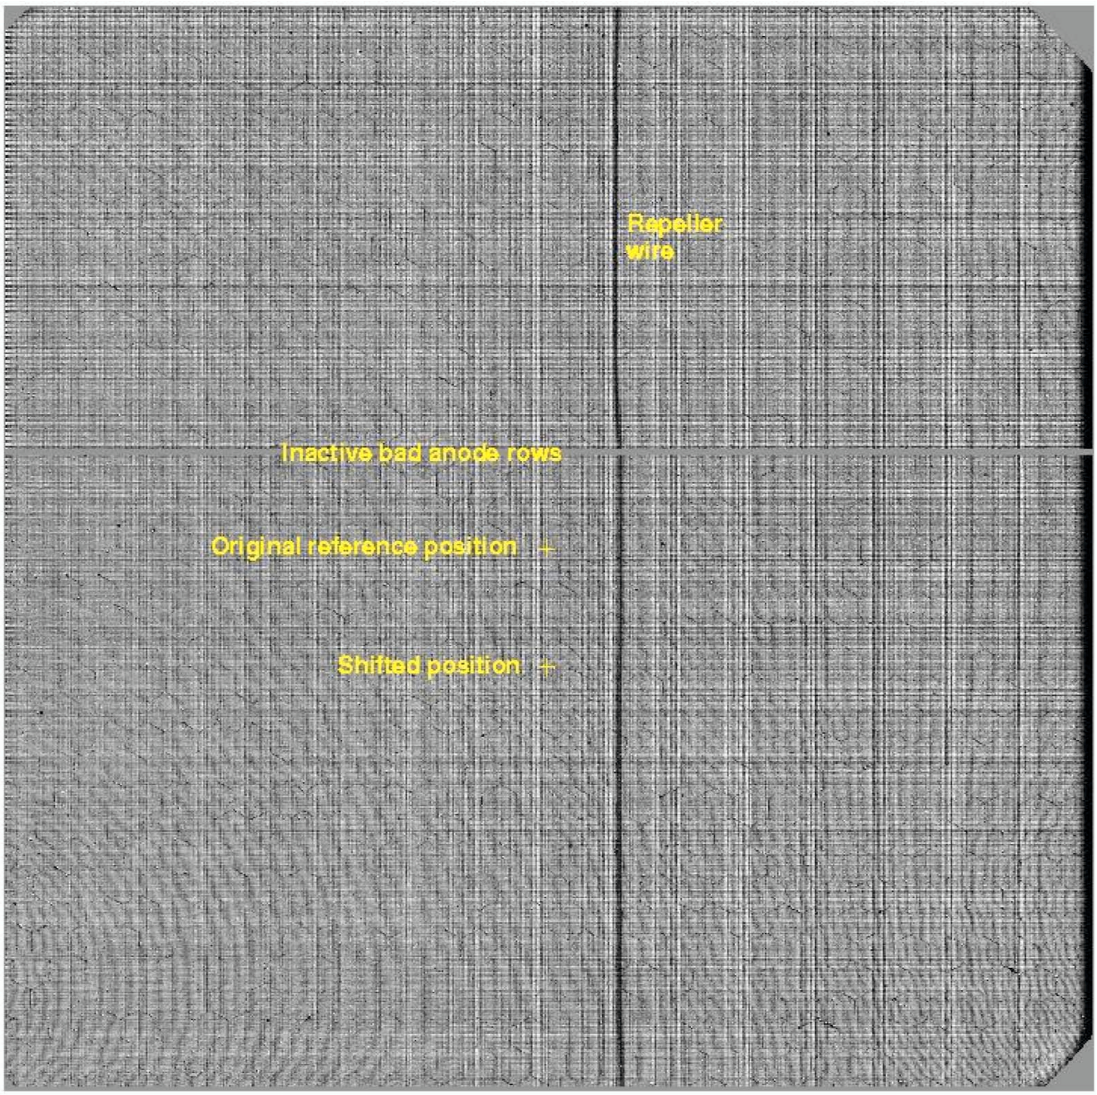

Figure 2: Features and problems on the SBC.

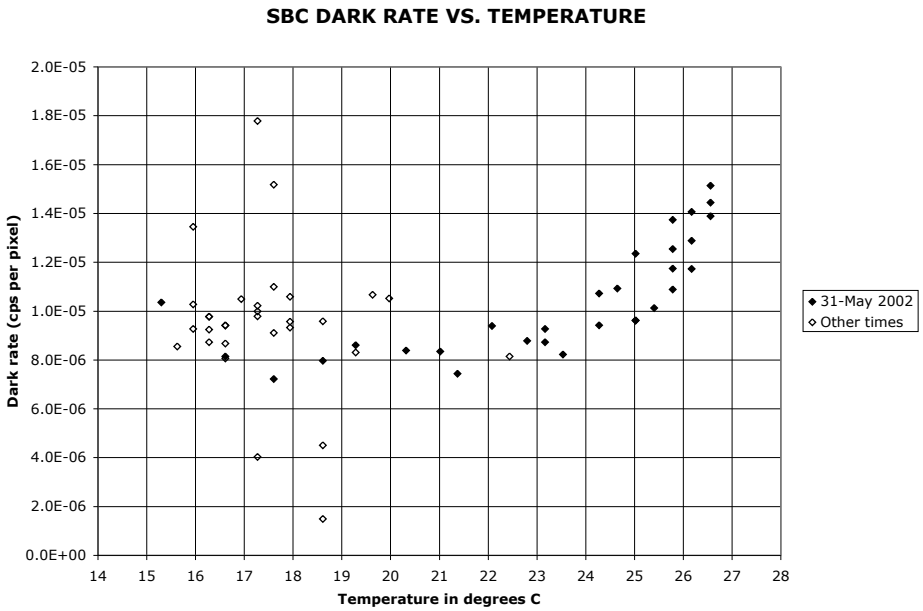

Figure 3: Variation of SBC dark rate with temperature.

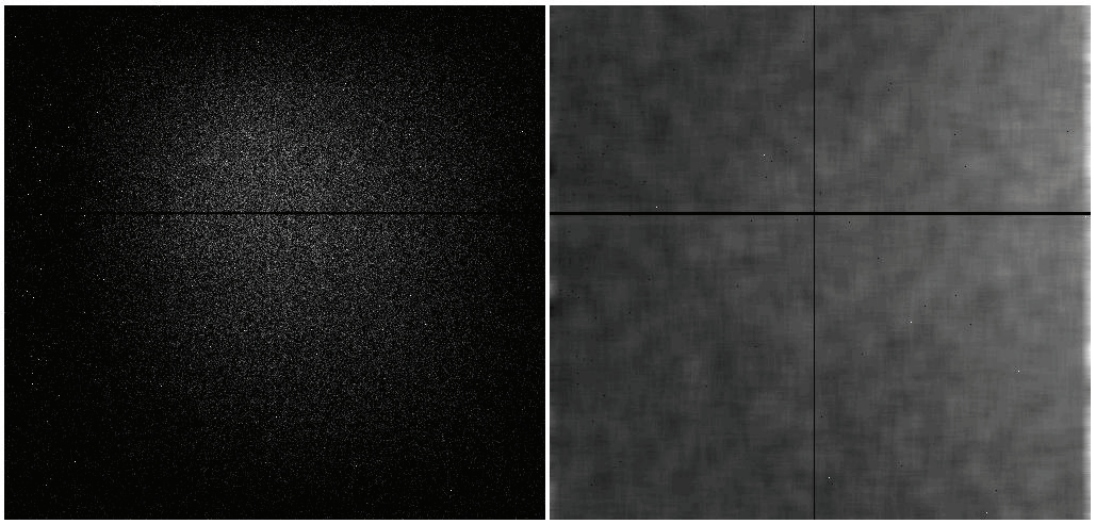

Figure 4: Old and new dark rate images.

expressed as polynomials which describe the overall distortion shape. For the WFC and HRC we have now incorporated a residual correction image which allows for just this sort of correction, going beyond the polynomial fit. This mechanism could well be used to include the intrinsic detector distortion once we have refined the SBC polynomial fit.

5. Spectroscopy

A report at this conference (Walsh, 2005) describes the calibration of the prism modes of the SBC and software to extract the spectra.

References

Tran, H. D. et al., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 86

Mack, J., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 67

Mack, J., Gilliland, R., van der Marel, R., & Bohlin, R. C., 2005, *Instrument Science Report* ACS 2005-13 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Bohlin, R. C., & Mack, J., 2005, *Instrument Science Report* ACS 2005-04 (Baltimore: STScI)

Walsh, J. R., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 79

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Recent Developments of the ACS Spectral Extraction Software aXe

Jeremy R. Walsh, Martin Kümmel, & Søren S. Larsen

Space Telescope – European Co-ordinating Facility, Karl-Schwarzschild-Strasse 2, D-85748 Garching, Germany

Abstract. The software package aXe provides comprehensive spectral extraction facilities for all the slitless modes of the ACS, covering the Wide Field Channel (WFC) grism, the High Resolution Channel (HRC) grism and prism and the Solar Blind Channel (SBC) prisms. The latest developments to the package apply to all ACS slitless modes leading to improved spectral extraction. Many thousands of spectra may be present on a single deep ACS WFC G800L image such that overlap of spectra is a significant nuisance. Two methods of estimating the contamination of any given spectrum by its near neighbors have been developed: One is based on the catalogue of objects on the direct image, another uses the flux information on multi-filter direct images. An improvement to the extracted spectra can also result from weighted extraction and the Horne optimal extraction algorithm has been implemented in aXe. A demonstrated improvement in signal-to-noise can be achieved. These new features are available in aXe-1.5 with the STSDAS 3.4 release.

1. The spectral extraction package aXe

As part of a collaborative project between STScI and the Advanced Camera for Surveys (ACS) IDT, the ST-ECF provides comprehensive support for the slitless spectroscopy modes of the ACS. As well as support to users of the ACS grism (G800L for the Wide Field Channel, WFC and High Resolution Channel, HRC) and prisms (PR200L for the HRC and PR110L and PR130L for the Solar Blind Channel, SBC) and contributions to the ground and in-orbit calibrations of the slitless modes, a primary pillar of this project has been the provision of an extraction package, called aXe. The package consists of a number of self-contained modules, which perform the basic steps - defining apertures for extraction of a spectrum, assigning wavelengths to pixels, flat fielding, extracting rectified 2D spectra and a 1D spectrum, and applying flux calibration. The modules are scripted in Python to allow easy integration into Pyraf/STSDAS (see Kümmel et al. 2005 for more details). The aXe user manual (`http://www.stecf.org/software/aXe/documentation.html`) provides full details for installing and running the package.

The fundamental aspect of slitless spectroscopy is that the individual objects define their own 'slit' in terms of position on the detector and object height of the dispersed spectrum; the object width in the dispersion direction affects the spectral resolution. aXe uses a catalogue of the observed targets, which is usually taken from a matched direct image as the starting point of the reduction process. In the design of aXe it was decided to make the software as general as possible, so that in the longer term not only slitless spectra from ACS could be extracted. This flexibility is engineered by putting all instrument specific parameters in a configuration file. Thus the specification of the spectral traces, the dispersion solutions, the name of the flat field file, the sensitivity file name, etc are all listed in a single file for each instrument mode. Thus for ACS, there are six configuration files; three for the G800L with the WFC (one for each chip) and HRC, one for PR200L, and one each for the SBC with PR110L and PR130L. The built-in flexibility has paid off since aXe

has also been used to extract spectra from multi-object spectra taken with the VLT FORS2 instrument (Kümmel et al. 2006).

Since the first release of aXe in 2002 ready for installation of ACS into *HST*, the package has evolved. In particular in 2004 a major enhancement was added with the use of 'drizzle' (Fruchter & Hook 2002) to combine 2D spectra of individual objects when the data is taken with dithers of the telescope. This turns out to be a very common observational procedure, at least for the grism modes, in order to recover the undersampling of the Point Spread Function (PSF) and to mitigate the effect of pixel sensitivity variations and hot pixels. Since the 2004 release (aXe-1.4), two enhancements have been added which are here described. The sensitivity of the ACS slitless spectroscopy modes implies that, despite the small pixel scale and compact PSF, even for high Galactic latitude fields, the surface density of detected spectra on moderate exposures (> thousands of seconds) displays crowding and overlap. A high priority was to indicate to the user which pixels in an extracted spectrum are affected by overlap with spectra of other objects. This has now been implemented in a quantitative way, whereby the estimated value of the contaminating flux contributing to a spectrum pixel is output. The second enhancement was to apply the well known technique of weighting by the spatial profile when forming a 1-D spectrum from the 2D spectrum on the detector (optimal extraction of Horne 1986). Both these enhancements are described.

2. Handling spectral contamination in aXe

The contamination of one spectrum by spectra of neighboring objects can be manifest in several ways:

- overlap of the first order spectrum by the first order spectrum of a nearby object situated in the projected direction of dispersion;
- overlap of first order spectra situated nearby but offset perpendicular to the dispersion direction;
- both above cases combined and possibly involving spectral orders other than the first.

A particular case is that of the zeroth order of a bright object overlying a fainter object spectrum. For the G800L grism, the zeroth order is similar in size to the dispersing object, but slightly dispersed, so that it can resemble a broad emission line. If such a feature is not recognized as a contaminating line it may lead to erroneous wavelength, and hence redshift, assignment. Of course the effect of bright objects on faint object spectra is more serious than the reverse case, and highlights the need for a warning of the contamination which provides at least an estimate of the actual contaminating flux to a given spectrum. In the first release of aXe, a minimal indicator of spectral contamination was implemented which indicated the total number of other object spectra which fell within a pixel in the extraction box of the given object. With aXe-1.5, a new scheme was implemented which provides a robust estimate of how much contaminating flux contributes to a given pixel. The contamination is estimated by making a simulated slitless image, which is achieved by one of two methods. The simpler method takes the input catalogue and generates simulated images as 2D Gaussians based on the image parameters, and is called the Gaussian Emission Model; the other method actually uses the fluxes in the pixels of a set of multicolor companion direct images, and is called the Fluxcube Emission Model.

2.1. Gaussian Emission model

The input catalogue which drives the object extraction is usually produced by running SExtractor on a companion direct image (or several images taken with different filters) and lists the object position, size parameters and magnitude. Thus for each object a spectrum

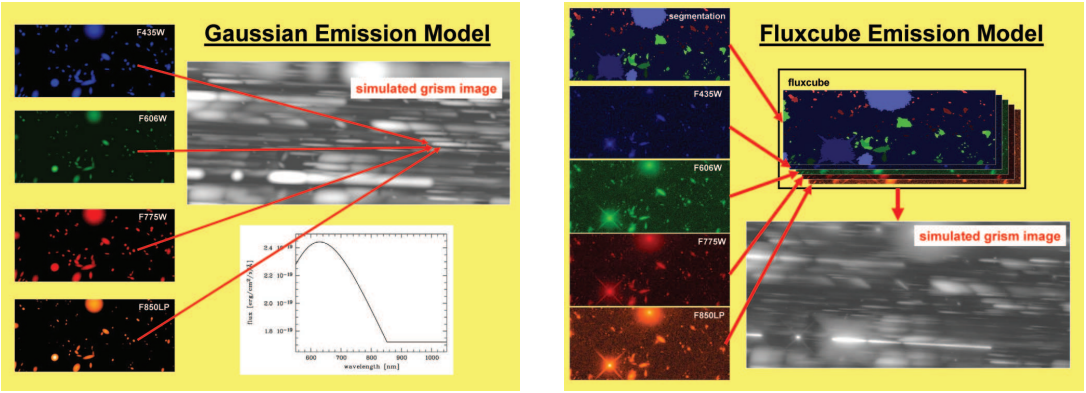

Figure 1: **Left.** The scheme for generating the contamination model based on 2-D Gaussians and an input catalogue with object shapes. The position, size (major and minor axes and position angle) and magnitude of objects in different filter direct images provide simulated objects which are dispersed onto the detector pixels, converted to e^- s^{-1}. **Right.** As left panel, but showing the procedure for generating the simulated images directly from the photometric information in the pixels of the multicolor direct images.

over the wavelength extent of the slitless dispersing element can be formed and converted to detector count rate using the known sensitivity of the slitless mode. For a single filter this will be a flat featureless spectrum, but more filter bands can more closely match the true object spectrum. From the position and extent of each object, which are determined from the object parameters of each image (A_IMAGE, B_IMAGE and THETA_IMAGE in SExtractor), a simulated spectrum corresponding to the slitless spectrum can be formed with spatial extent matching the object size (see Figure 1, left panel, for an overview).

2.2. Fluxcube Emission model

An alternative, and preferable, method to produce the model spectrum for estimation of the contribution of contaminating spectra is to use the surface brightness distribution in the pixels of the companion direct image(s). The assignment of pixels to a given object still has to be established and the method chosen in aXe-1.5 was to use the Segmentation image provided by SExtractor. This image has the pixel value belonging to a given object set to the object number in the SExtractor catalogue. The data are stored in an intermediate file which is a cube with planes for the segmentation image and the filter images in different bands; it is hence called a Fluxcube (see Figure 1, right panel). aXe-1.5 provides a routine to produce such flux cubes which are then read by the appropriate tasks during the extraction of the spectra.

2.3. Quantitative contamination

During the extraction process for a given object spectrum, then, within the extraction box, the count rate in the pixels belonging to objects other than the one being extracted are accumulated from the contamination image (originating in either the Gaussian or Fluxcube models) and assigned to the contaminating signal. Applying the sensitivity curve to the contaminating signal enables the total contaminating signal to the spectrum to be computed. The result is written to a separate column of the Extracted Spectra File by aXe. Figure 2 shows as an example the extracted spectrum of an $i \sim 24$ mag. emission line galaxy; the strong red continuum is shown to be badly contaminated (squares show the contamination contribution). However the pair of emission lines at around 7500 Å can be seen to be intrinsic to the object and not arise from a contaminating spectrum. It must be emphasized that the quantitative contamination is only an estimate as it is based on a model (Gaussian

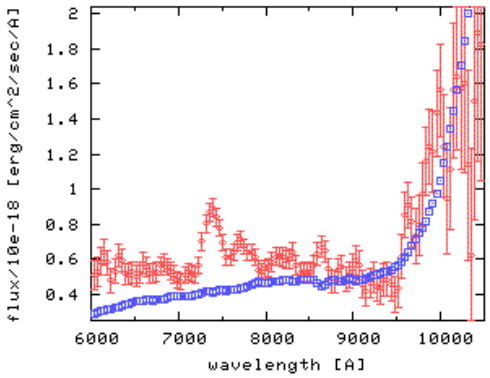

Figure 2: The extracted ACS WFC 1D spectrum of an emission line galaxy from the HUDF parallel data (primary instrument was NICMOS (PI: Thompson, programme 9803). The flux of the galaxy (with 1σ error bars) (diamonds) and the sum of the contaminating flux in the extraction box (boxes) are shown.

or Fluxcube); it does however lead to an appropriate level of caution being exercised in quantitatively assessing a given spectrum.

The aXe tasks for producing the emission model leading to contamination estimation deliver simulated images from direct image(s) or a catalogue based on direct image(s). Thus the slitless images produced by the Gaussian or Fluxcube models (....CONT.fits files) are ideally suited to use in observation planning and exposure time estimation. For example for a complex field, the contamination images can be used to choose optimal telescope roll angles that minimize overlap of the spectra of interest.

3. Weighted spectral extraction

If the 2D spectrum of an object is extracted applying a set of weights to the spatial profile, then the resulting 1D spectrum has a higher signal-to-noise than the simple box-extracted (i.e. summed) spectrum. This was shown by Horne (1986; see also Robertson 1986) and is often referred to as optimal extraction. Weighted extraction has been implemented in aXe-1.5, with the Horne (1986) algorithm, using weights derived from the contamination image described in Section 2. The contamination image is well suited to providing the weights since it is binned in the same way as the spectrum to be extracted and, as a model, does not suffer from any systematic statistical effects. In the examples described by Horne (1986), the weights are typically determined by fitting the observed spectrum in the dispersion direction as a function of spatial offset. This procedure can be prone to failure for weak spectra and where the whole spectral extent is not occupied by signal; in the case of overlapping spectra it would provide incorrect weights.

Simulations were performed with a random star field composed of (spatially) well-sampled star images in order to test the optimal extraction implementation in aXe-1.5. The quantitative contamination procedure in aXe using the Gaussian emission model was employed to make simulated slitless spectra from a SExtractor catalogue; background and noise were added and then the spectra were extracted with and without weighting. The gain of weighted extraction was specified as the ratio of the signal-to-noises in the optimal to the unweighted spectrum over a given range. As the signal-to-noise level decreases, the optimal extraction shows an advantage over the unweighted extraction depending on the width of the extraction (see Figure 3 where the extraction width is 6 times the object width). At the lowest signal-to-noise level, the advantage is around a factor 1.4, equivalent to an increase in exposure time of 1.9 over unweighted extraction.

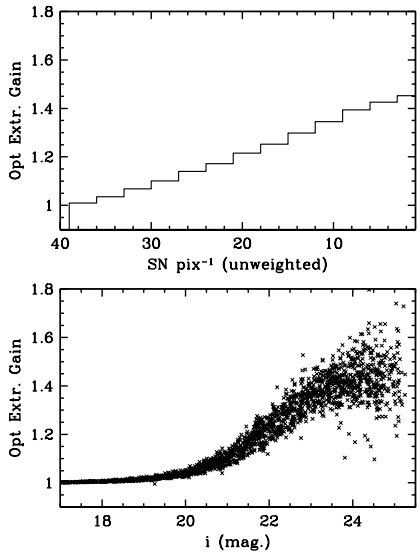

Figure 3: The result of a simulation of a star field observed with the ACS WFC G800L grism in terms of the signal-to-noise (S/N) advantage of the weighted over the unweighted extraction (extraction widths ± 3 σ of the sources). The lower plot shows the actual advantage for each star in the simulation (ratio of weighted S/N divided by unweighted S/N, averaged over a wavelength range) in terms of the i mag. The upper plot shows a histogram version of the lower plot where the plot abscissa is the mean signal-to-noise over a 1000Å range.

As an example of weighted extraction applied to real ACS data, a crowded stellar field was selected from archive data resulting from the APPLES programme (ACS Pure Parallel Lyman-alpha Emission Survey, PI: J. Rhoads, Programme 9482). Over 7000 spectra (to $i =$ 25 mag) were present on this ACS Wide Field Channel image so there is often considerable contamination between adjacent spectra. Using the quantitative contamination, only those spectra with <5% contamination were selected, and the advantage in signal-to-noise of using weighted extraction was determined. The mean S/N was computed over a range of 1000 Å centred at 7500 Å and Figure 4 shows the result as a point plot v. magnitude and as a histogram v. signal-to-noise. Here 1374 stars are analyzed, and a peak advantage of weighted over unweighted extraction of about 1.3 is realized. This is somewhat lower than that achieved in the simulations, but is probably on account of the narrow undersampled Point Spread Function (PSF) of the WFC data. Nevertheless the theoretical gain in exposure time is around 60% at low signal-to-noise.

Weighted extraction can be selected in the aXe parameter set and both weighted and unweighted 1D spectra for all sources are output. For sources with complex cross dispersion profiles, there will probably be little advantage of weighted extraction but for small objects, such as faint stars and distant galaxies, a modest gain in signal-to-noise is achievable with optimal extraction for slitless spectra.

The latest release aXe-1.5 provides both quantitative contamination and weighted extraction for all ACS slitless spectroscopy modes and is available from November 2005, with Pyraf/STSDAS 3.4. Full details can be found at
`http://www.stecf.org/software/aXe/`

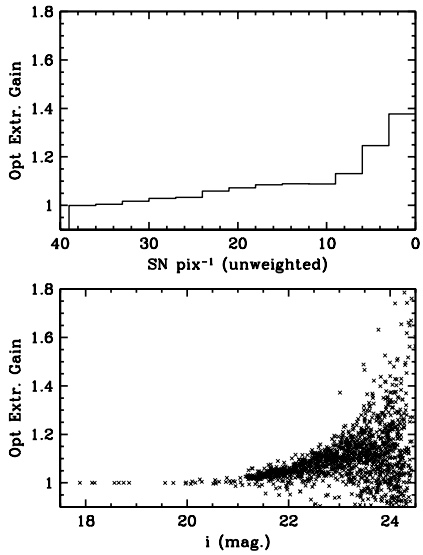

Figure 4: An identical plot to Figure 3, but now based on real data taken with the ACS WFC G800L grism as part of *HST* APPLES programme 9482. The field is at low Galactic latitude and consists of many late type stars. There is considerable scatter in the weighted v. unweighted extraction advantage at the lowest signal level, but a demonstrable increase in signal-to-noise at low flux levels is achieved.

References

Horne, K., 1986, PASP, 98, 609

Fruchter, A. S., & Hook, R. N., 2002, PASP, 114, 144

Kümmel, M., Larsen, S. S., & Walsh, J. R., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 85

Kümmel, M., Kuntschner, H., Larsen, S. S, & Walsh, J. R., 2006, in Astronomical Data Analysis Software and Systems XV. Eds., C. Gabriel, C. Arviset, D. Ponz, E. Solano (Astronomical Society of the Pacific: San Francisco), in press

Robertson, J. G., 1986. PASP, 98, 1220

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Slitless Spectroscopy with the Advanced Camera for Surveys

M. Kümmel, S. S. Larsen, and J. R. Walsh

Space Telescope – European Coordinating Facility, Karl-Schwarzschild-Str. 2, D-85748 Garching b. München, Germany

Abstract. The Advanced Camera for Surveys (ACS) enables slitless, low-resolution spectroscopic imaging in its three channels. The most-used modes are grism imaging with the WFC and the HRC at a resolution of 40 and 24 Å/pixel, respectively. In the far-UV there are two prisms for the SBC and one prism for the HRC in the near-UV. An overview of the slitless spectroscopic modes of the ACS is presented together with the advantages of slitless spectroscopy from space. The methods and strategies developed to establish and maintain the wavelength and flux calibration for the different channels are outlined. Since many slitless spectra are recorded on one deep exposure, pipeline science quality extraction of spectra is a necessity. To reduce ACS slitless data, the aXe spectral extraction software has been developed at the ST-ECF. aXe was designed to extract large numbers of ACS slitless spectra in an unsupervised way based on an input catalogue derived from a companion direct image. In order to handle dithered slitless spectra, drizzle, well-known in the imaging domain, has been applied. For ACS grism images, the aXedrizzle technique resamples 2D spectra from individual images to deep, rectified images before extracting the 1D spectra. aXe also provides tools for visual assessment of the extracted spectra and examples are presented.

1. Grisms and prisms on the ACS

The ACS has three channels, the Wide Field Channel (WFC), the High resolution Channel (HRC) and the Solar Blind Channel (SBC), and each channel is capable of delivering slitless spectroscopic images by inserting a grism or prism into the optical beam. The five different combinations of ACS channel and dispersing element offer low resolution spectroscopy from the UV to the far-red wavelength regime. Table 1 lists important parameters such as spectral resolution, wavelength range and field of view (FoV) for all slitless modes of ACS.

Table 1: The spectroscopic modes of the ACS.

Channel	Disperser	Wav. Range [Å]	Resolution [Å/pixel]	Pixel Size [mas/pixel]	FOV [arcsecond]
WFC	G800L	$5500 - 10500$	38.5	50×50	202×202
HRC	G800L	$5500 - 10500$	23.5	28×24	29×26
HRC	PR200L	$1700 - 3900$	20[@2500Å]	28×24	29×26
SBC	PR130L	$1250 - 1800$	7[@1500Å]	34×30	35×31
HRC	PR110L	$1150 - 1800$	10[@1500Å]	34×30	35×31

Figure 1 illustrates the differences of using grisms vs. prisms as dispersive elements by comparing the dispersion and the sensitivity of the HRC grism G800L (left) and the HRC prism PR200L (right). The G800L has an almost constant dispersion over the entire accessible wavelength range. The dispersion of the PR200L increases drastically towards longer

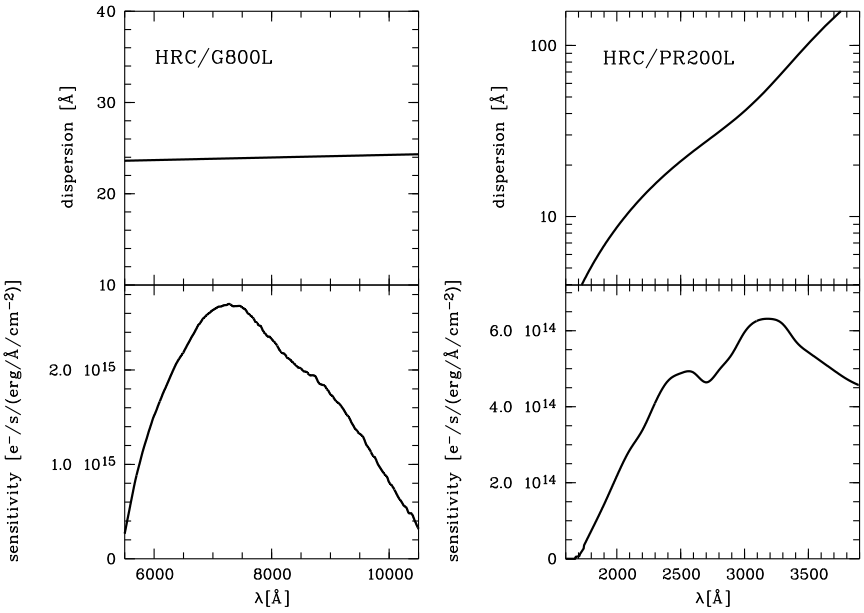

Figure 1: A comparison between dispersion and sensitivity of the HRC grism G800L (left) and prism PR200L (right).

wavelengths, and any value for the dispersion must be accompanied by the wavelength at which it is specified (see Tab. 1 and the ACS Instrument Handbook). The different properties of the prisms and grisms require a flexible data reduction software to be able to reduce both slitless grism and prism data.

Following the demise of STIS in August 2004 the interest in using ACS grisms and prisms has increased substantially, since it is now the only optical-UV spectral capability aboard HST. As a result, around 10% of all approved orbits in Cycle 14 are devoted to slitless spectroscopy mode with ACS (Macchetto et al. 2005).

2. Slitless spectroscopy from space

Slitless spectroscopy with the ACS has some distinctive features which clearly separate the technique from spectroscopy with slits. Its advantages are:

- The sky background is extremely low. Typical background count rates for the combinations WFC/G800L, HRC/G800L and HRC/PR200L are 0.1, 0.006 and 0.04 e^-/s/pix, respectively. With typical exposure times in the range $1000-2000$ s the read-out noise of $\sim 4.9\,e^-$ remains an important to even dominant source to the overall noise in the images.

- In contrast to the ground-based sky background, which is usually dominated by emission lines, the HST sky background is much smoother, which makes the background removal less problematic. Since the removal of a background with emission lines always leaves variations in the signal-to-noise ratio and therefore leads selection effect in the analysis of the data. ACS slitless spectroscopy avoids these selection effects.

- ACS slitless spectroscopy is associated with an extremely large data yield. To illustrate this, Figure 2 displays a MultiDrizzle-combined WFC/G800L image. While it is not

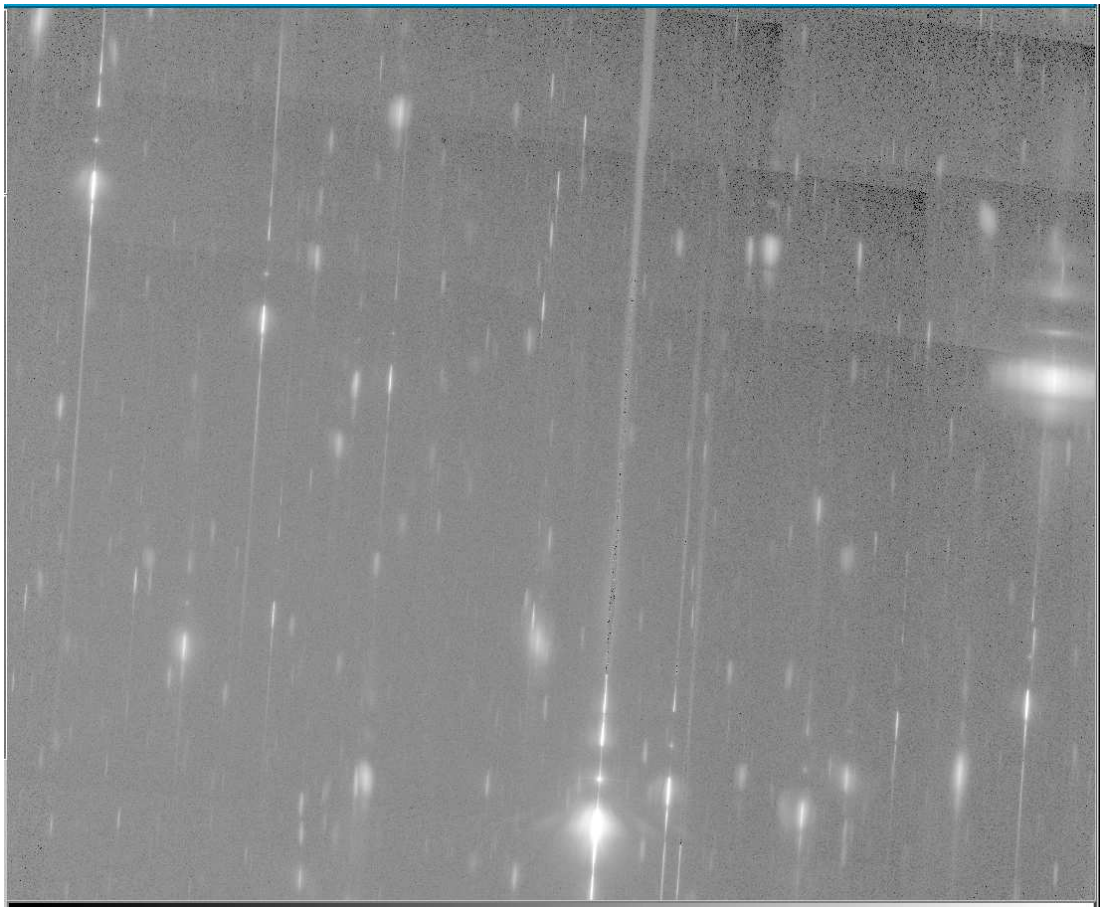

Figure 2: A MultiDrizzle-combined WFC/G800L image. The large number of spectra to be extracted from the data and the contamination of spectral orders are evident.

general to extract slitless spectra from MultiDrizzle-combined images such as Fig. 2, such co-added images give a first impression on what can be expected from the data. The image shows the spectra of thousands of objects, which all can be extracted from the data. In ACS slitless spectroscopy, the number of objects per field to be extracted is solely determined by the depth of the observation.

- Also the typical advantages of the HST, such as the compact point spread function and the high stability of the instrument, apply to ACS slitless spectroscopy.

The disadvantages of the ACS slitless spectroscopy modes are:

- The contamination of spectra is an ubiquitous phenomenon in slitless spectroscopy, especially since the absence of slits or masks allows a contamination in the spectral direction (see Fig. 2). The contamination can occur even from higher spectral orders. In Fig. 2 this becomes evident when looking at the various spectral orders of the bright stars, which cover a large area and of course overlay the spectra of the fainter objects. Contamination affects all spectra to various degrees, and the reduction software has to deal with contamination (see Section 5 and Walsh, Kümmel & Larsen 2005).

- The reduction of slitless spectroscopic data is quite different from the usual extraction of spectra taken with slits. It requires different methods, the application of different

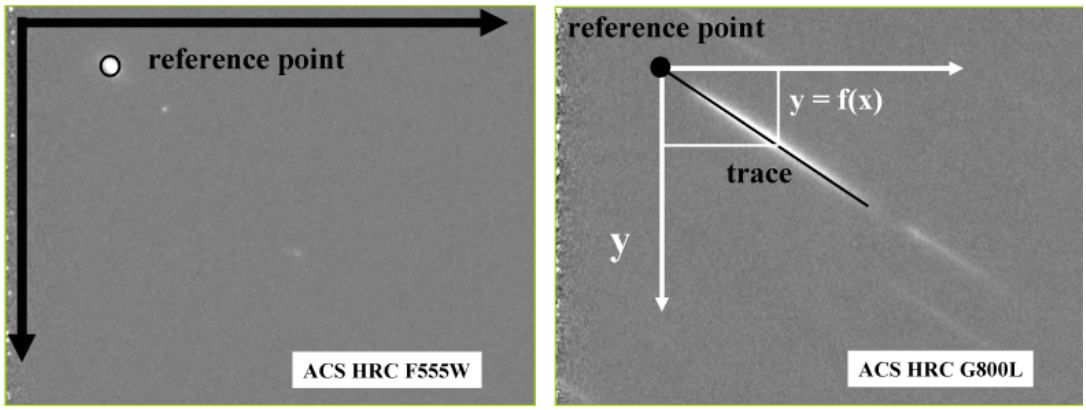

Figure 3: The HRC/G800L grism image (right panel) and the associated direct image taken with the F555W filter (left panel). The reference point for an example object is determined with SExtractor in the F555W image and then transferred to the G800L image. In the grism image the reference point is the origin of the coordinate system to describe the location of the trace. Finally, the wavelength solution is applied as a function of the trace length.

calibrations and the usage of different software. As a consequence, the astronomer needs some time to get familiar with slitless data and its reduction concepts.

3. Reducing slitless spectroscopic data

In 'traditional' spectroscopy with slits, the aperture (slit or mask) defines a framework for the trace definition and the wavelength solution. The information derived from calibration data (flat field- and arc-exposures) taken with the identical instrumental setup is directly transferred to the science data to extract the spectra.

In slitless spectroscopy, however, such a framework does not exist. The exact location of the objects on the science data is unknown a priori, and it is impossible to establish a trace description and a wavelength solution on the basis of slitless data alone such as in the right panel in Figure 3. To base a framework or coordinate system for the trace description and wavelength solution, a so-called *reference point* is needed for every spectrum which is to be extracted. The reference point is the basis for the origin of a local coordinate system (see Fig. 3, right panel) to define the trace description and apply the wavelength solution. In ACS slitless spectroscopy, the reference points of all objects are established on a direct image such as shown in Fig. 3, left panel, which is associated with every slitless image. Then the extraction procedure is as follows:

1. All object positions and therefore reference points are determined on the direct image, which was taken very close in time to the slitless image. This is done either directly with a source detection software such as SExtractor, or indirectly by computing their pixel coordinates on the direct image (left panel in Fig. 3).

2. The reference positions are transferred to the slitless image (right panel in Fig. 3).

3. A so-called *aXe configuration file*, which was assembled using calibration data (see next Section), gives the prescription to define the trace description and the wavelength solution for every reference position on the slitless image.

4. The object spectra are finally extracted from the slitless image.

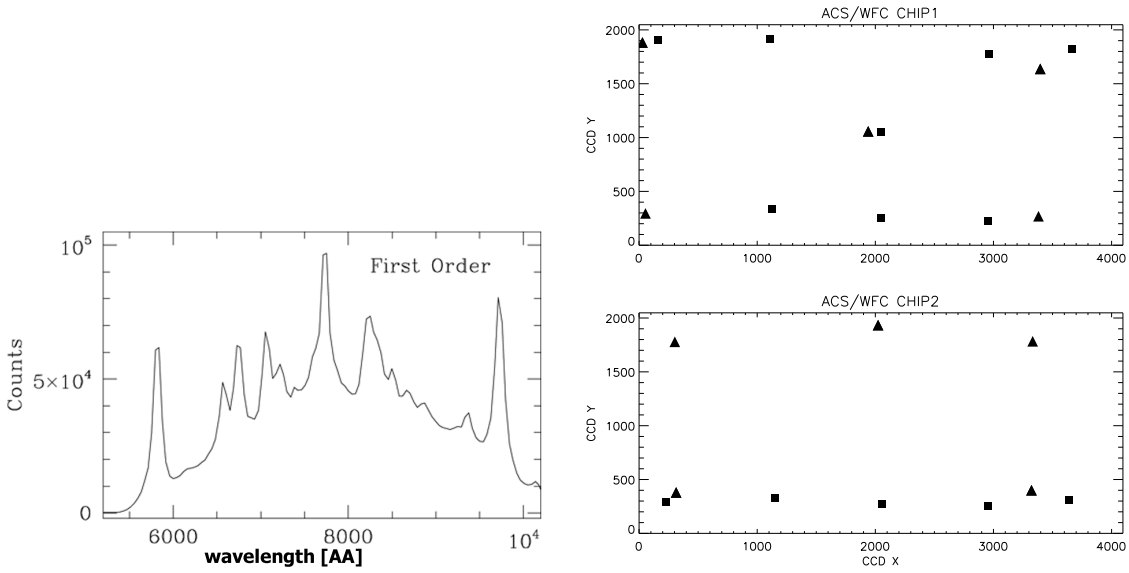

Figure 4: Left panel: The spectrum of the Wolf-Rayet star W45, which was used to establish a wavelength calibration for the WFC/G800L mode. Right panel: The positions of astronomical targets across the WFC/G800L FoV to determine the 2D-dependence of the calibrations.

4. The calibration of the slitless modes

The main calibration products acquired for an ACS slitless mode are:

- The aXe configuration file to describe trace description and wavelength solution.
- A sensitivity file for the flux calibration of every spectral order.
- A flat field to be able to apply pixel-to-pixel gain corrections.

Since there are no on-board calibration lamps for ACS, suitable astronomical targets with known fluxes and known spectral features were observed in dedicated calibration proposals to establish the flux calibration and wavelength calibration, respectively.

For the flux calibration and trace definition, white dwarf standards could be used for all ACS slitless modes. For the wavelength calibration, planetary nebulae and Wolf-Rayet stars were observed for the optical modes WFC/G800L and HRC/G800L (Pasquali et al. 2005, Larsen & Walsh 2005). The left panel of Figure 4 shows the first-order spectrum of the Wolf-Rayet star WR45 as observed with the WFC/G800L. The bright emission lines which spread over the entire spectral range are identified to define the wavelength solution. The ACS prism modes, which cover the UV (HRC/PR200L, SBC/PR120L and HRC/PR110L), were wavelength calibrated using observations of planetary nebulae and carefully redshift-selected QSO's (Larsen, Kümmel & Walsh 2005).

In all ACS slitless modes, the trace definition, wavelength calibration and sensitivity show variations across the FoV. For this reason the astronomical calibration targets were observed at several positions across the FoV. Figure 4 shows as an example the positions of the astronomical calibration targets across the WFC FoV. After establishing the calibration at each position individually, a global 2D solution is fitted to the individual solutions, which then enables one to derive the calibration at every arbitrary position in the FoV.

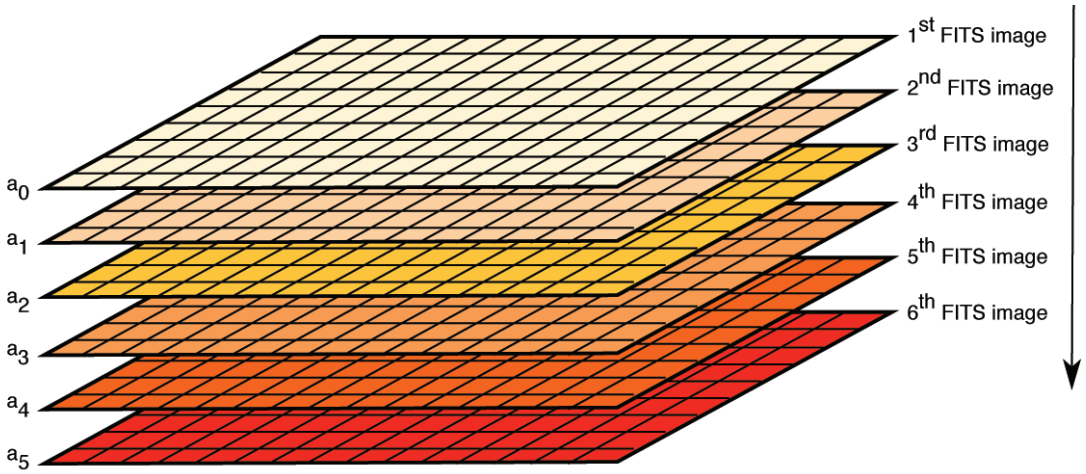

Figure 5: The structure of a flat field cube. Each extension of the FITS file gives the coefficients to compute the flat field correction for every pixel.

The absence of on-board calibration lamps also inhibits the conventional approach for flat fielding. But even if calibration lamps were present, taking a flat field exposure for every science exposure (as is usually done for slitted spectroscopy), could not deliver a proper flat field. In ACS slitless spectroscopy the objects and their exact position, and therefore the wavelength attributed to the pixels, are unknown prior to the data reduction. Moreover there are several wavelengths associated to each pixel in regions where the spectra of multiple objects overlap, and this is not possible in the conventional approach.

As a consequence, the flat field used in ACS slitless spectroscopy must be able to correct any pixel for any wavelength. The solution to this problem is a flat field cube, as shown in Figure 5. A flat field cube is a multi-extension FITS file, and every extension contains the coefficients to compute the gain correction for any given pixel (i, j) at any wavelength λ according to the formula:

$$FF(i, j, \lambda) = a_0(i, j) + a_1(i, j) * \lambda + a_2(i, j) * \lambda^2 + a_3(i, j) * \lambda^3 \ldots$$

The flat field cubes for the different spectroscopic modes are created from photometric flats, and further details on the flat fielding is given in Walsh & Pirzkal (2005).

5. The aXe spectral extraction software

The aXe software package was specifically designed to extract spectra in a consistent manner from all slitless spectroscopic modes of the ACS. While the reduction of slitless spectroscopic ACS data was the driving force behind the development of aXe, the software package was designed to allow also the reduction of spectroscopic data from other instruments without the need for fundamental software changes, and currently an investigation is done to apply aXe to data taken with the FORS2-MXU unit at the VLT (see Küemmel et al. 2006 and Kuntschner et al. 2005).

aXe is a PyRAF/IRAF package with several tasks (see Figure 6) which can successively be used to produce extracted spectra from slitless data. As can be seen in Fig. 6, there exist two classes of aXe tasks:

1. The Low Level Tasks work on individual images. Their aim is to perform a certain reduction step on a particular image.

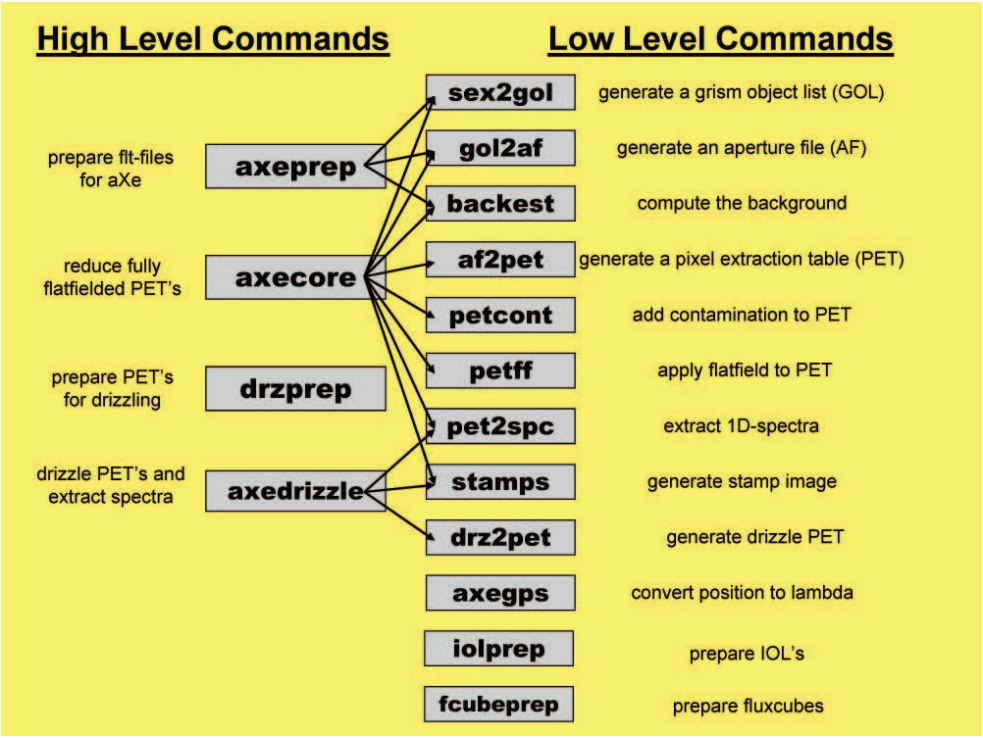

Figure 6: The list of aXe tasks. The arrows indicate the interaction between the High Level and the Low Level Tasks. The right column describes the reduction step executed by a certain task.

2. The High Level Tasks work on data sets. Their aim is to do a sequence of processing steps on several images.

Often the High Level Tasks call Low Level Tasks to perform a certain reduction step on individual images, and Fig. 6 shows their interaction. The High Level Tasks were designed to cover all steps of the aXe data reduction, and working with aXe usually means to apply the High Level Tasks to data sets. Due to the large data yield in ACS slitless data, which makes it impossible to extract every spectrum individually, aXe is built as a semi-automatic data reduction system. After the object positions are determined on the direct images (see Sect. 3.), aXe runs automatically and creates (besides the final, extracted spectra) additional intermediate products, such as 2-dimensional stamp images of the spectra. The intermediate data products help the user to check the reduction procedure and to fine-tune the extraction parameters.

The aXe tasks are implemented in Python. To work on the pixel data (such as flat fielding or extracting 1D spectra), which requires a large computational speed, the Low Level Tasks call executables which are implemented in standard **C**. Via PyRAF the aXe tasks are fully embedded into the STSDAS software package, and aXe users do not have to leave their familiar data reduction environment to work with ACS slitless data. The aXe package evolves continuously, and together with STSDAS new aXe versions are released about once a year. The current aXe-1.5 was released together with STSDAS 3.4 in November 2005. Between these large releases there are occasional smaller software releases via the aXe web page (http://www.stecf.org/software/aXe/) to publish bug-fixes or together with new aXe configuration files.

The two main drivers behind the development of the aXe software package are:

1. Improvements for the user convenience.

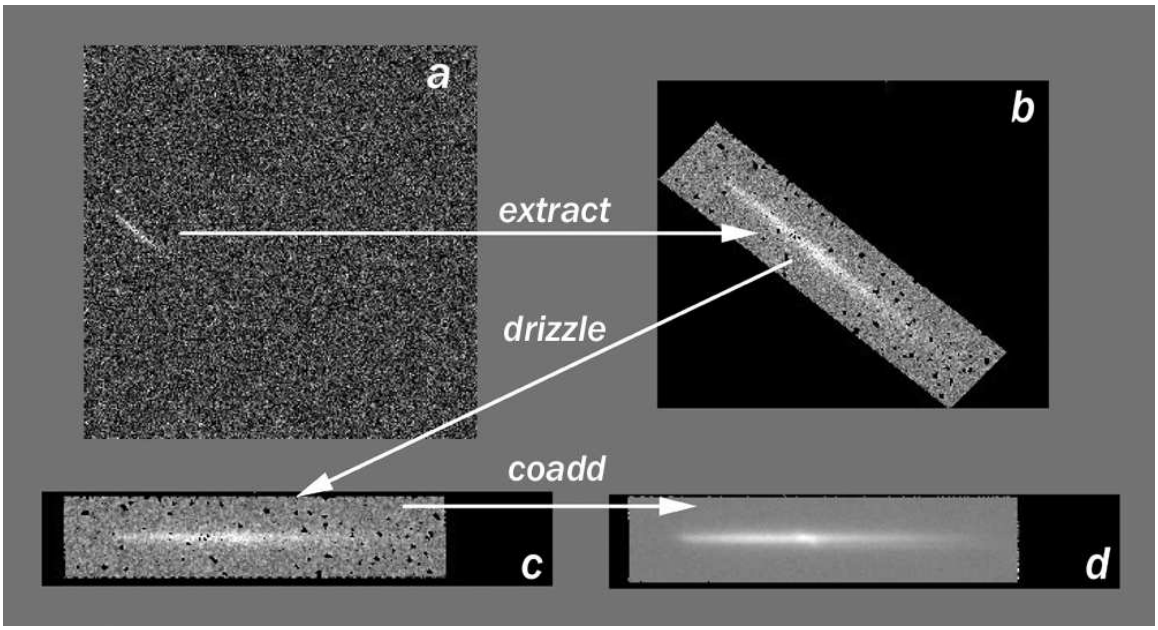

Figure 7: The aXedrizzle method of combining 2D spectra. The object in panel a is extracted as a stamp image (b), which is drizzled to an image with constant dispersion and constant pixel scale in spatial direction. The 1D spectrum is finally extracted from the deep, co-added 2D spectrum (d).

2. Adding new functionality to aXe.

The former was the motivation for developing the High Level Tasks in aXe-1.4, which greatly reduced the number of commands with which the user has to become acquainted. In aXe-1.5 the task `iolprep` was added, which is a new tool to generate SExtractor lists with object positions for a set of slitless images in a standard scenario.

As new functionalities, *optimal weighting* and the so-called *quantitative contamination* were added in aXe-1.5. The old aXe method of recording contamination associates to every pixel in the slitless image the number of spectra of which the pixel is a part. This information is processed in the 1D spectral extraction, and each spectral element has as contamination information the number of other spectra its input pixels contributed to. This method of recording contamination gives no clue on how much the contaminating objects actually influence the extracted object flux.

In quantitative contamination the contributing flux from other objects to every spectral bin is estimated according to a simple emission model. The quantitative contamination estimate is a very good tool for the user to decide which data points he can trust. A detailed description of the quantitative contamination method and of the optimal weighting implemented in aXe-1.5 is given in Walsh, Kümmel & Larsen (2005).

Another important addition to the functionality of aXe is the *axedrizzle* method, which was first released in aXe-1.4 (Kümmel et al. 2005). With the aXedrizzle method the individual 2D spectra of an object (see Figure 7a, b) are coadded to a single, deep 2D image (Fig. 7d) before extracting the 1D spectrum. The combination of the individual 2D spectra is done with the "Drizzle" software (Fruchter & Hook 2002), which is well known from HST imaging. This method of combining the data has several advantages:

- Resampling to a uniform wavelength scale and an orthogonal spatial direction with a constant pixel scale is achieved in a single step.
- Pixel weighting is handled correctly.

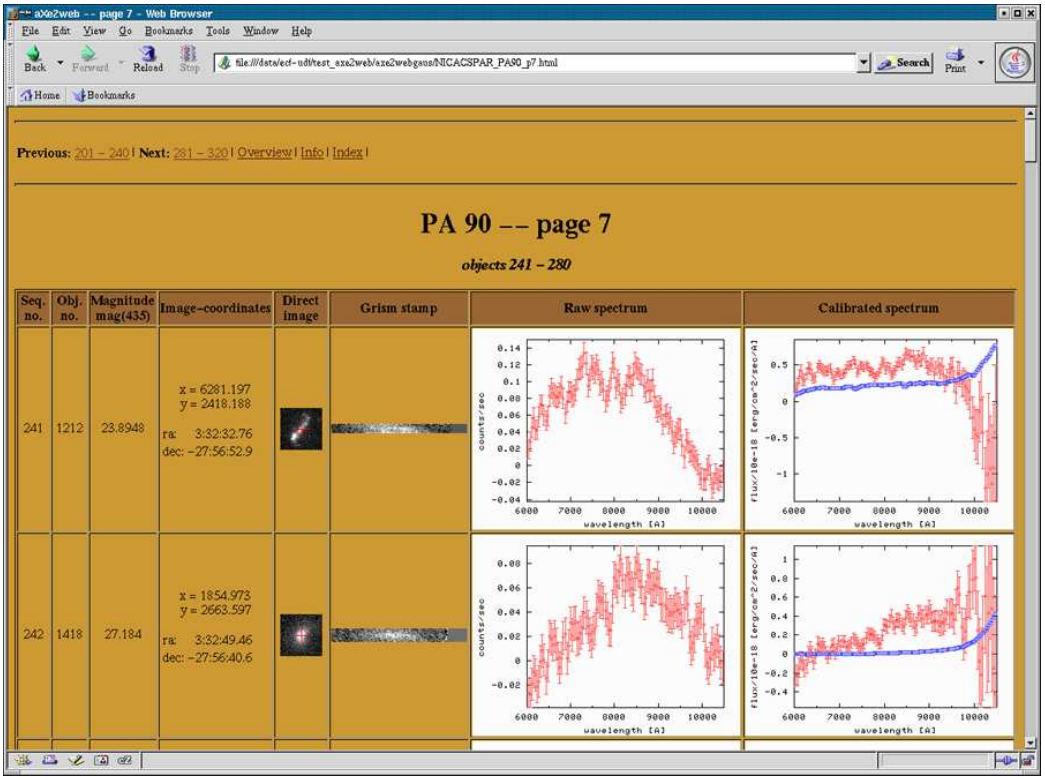

Figure 8: A web page created by aXe2web. The various rows contain information on individual objects such as their reference number, magnitude, position, stamp images and spectra

- The coadded 2D spectrum can be viewed to detect any problems.

Due to the highly non-linear form of the prism dispersion (see Fig. 1), the aXedrizzle method is restricted to the ACS grism modes WFC/G800L and HRC/G800L only.

6. The aXe visualization of spectra

As indicated above, a data set with ACS slitless images may contain hundreds or even thousands of spectra, and a visual inspection of each individual spectrum is very tedious. To help the user digest the large amount of data, the tool aXe2web was implemented, which produces browsable web pages for a fast and discerning examination of spectra (Walsh & Kümmel 2004).

aXe2web uses the direct image, the SExtractor catalogue, the aXe stamp images and the extracted spectra to produce an HTML summary containing a variety of information for each spectrum. Figure 8 shows part of an HTML page produced by aXe2web. Each object produces a line in the HTML page which lists the sequence number, the reference number, the magnitude of the direct image object, the coordinates (image and RA/DEC) of the direct image object, the direct stamp image, the grism/prism stamp image and the 1D extracted spectrum in *counts/s* and flux.

To facilitate an easy navigation within a data set, an overview and an index page accompany the object pages (Fig. 8) which show the detailed object information. The overview page contains for each object the basic information, e.g. object positions and magnitudes, and the index page includes a table with the ordered reference numbers of all

objects. Direct links guide from the overview page and the index page to the corresponding location on the object pages.

7. Conclusions

- All aspects of ACS slitless spectroscopy (calibration, software etc.,) are supported such that users can obtain and reduce slitless data in a pipeline way.
- ACS slitless spectroscopy is successfully used in various science projects such as the HUDF HRC Parallels (Walsh, Kümmel & Larsen 2004), high redshift supernovae research (Riess et al. 2004) and the search for high redshift galaxies (the GRAPES and PEARS programs, see Pirzkal et al., 2004).
- More information about ACS slitless spectroscopy, the calibration and the aXe software is given on the aXe web pages at http://www.stecf.org/software/aXe/
- User support concerning all topics related to ACS slitless spectroscopy is provided by the ACS group at the Space Telescope - European Coordinating Facility (ST-ECF). The centralized email address for requests is acsdesk@eso.org.

References

Fruchter, A. S., & Hook, R. N., 2002, PASP, 114, 144

Kümmel, M., Kuntschner, H., Larsen, S. S., & Walsh, J. R., 2006, in Astronomical Data Analysis Software and Systems XV, eds. C. Gabriel, C. Arviset, D. Ponz & E. Solano, ASP Conference Series

Kümmel, M., Walsh, J. R., Larsen, S. S., Hook, R. N., 2005, in Astronomical Data Analysis Software and Systems XIV, eds. P. Shopbell, M. Britton & R. Ebert, ASP Conference Series

Kümmel, M., Walsh, J. R., & Larsen, S.S., 2004, ST-ECF Newsletter 37, p. 14

Kuntschner, H., Kümmel, M., Larsen, S. S, & Walsh, J. R., 2005, ESO Messenger 122

Larsen, S. S., Kümmel, M., & Walsh, J. R., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 103

Larsen, S. S., & Walsh, J. R., 2005, *Instrument Science Report* ACS 05-08 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Macchetto, D., Williams, R., & Blacker, B., 2005, STSCI-Newsletter 22, p. 5

Pasquali, A., Pirzkal, N., Larsen, S. S., Walsh, J. R., Kümmel, M., 2005, PASP, accepted (astro-ph/0510428)

Pirzkal, N., Xu, C., Malhotra, S., et al., 2004, ApJS 154, 501

Riess, A., Strolger, L. G., Tonry, J., et al., 2004, ApJ 607, 665

Walsh, J. R., & Kümmel, M., 2004, ST-ECF Newsletter 35, p. 9

Walsh, J. R., & Pirzkal, N., 2005, ACS ISR 05-02

Walsh, J. R., Kümmel, M.,& Larsen, S. S., 2004, ST-ECF Newsletter 37, p. 8

Walsh, J. R., Kümmel, M.,& Larsen, S. S., 2006, this volume, 79

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Selection and Characterization of Interesting Grism Spectra

Gerhardt R. Meurer
Department of Physics and Astronomy, The Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218

Abstract. Observations with the ACS Wide Field Camera and G800L grism can produce thousands of spectra within a single WFC field producing a potentially rich treasure trove of information. However, the data are complicated to deal with. Here we describe algorithms to find and characterize spectra of emission-line galaxies and supernovae using tools we have developed in conjunction with off-the-shelf software.

1. Introduction

The G800L grism combined with ACS's Wide Field Camera is a powerful combination for obtaining thousands of spectra with relatively modest outlay of *HST* time. However, the resulting images are difficult to interpret due to a number of peculiarities including: (1) strong spatially varying sky background; (2) a position-dependent wavelength solution; (3) the wide spectral response: a three-dimensional flat field and modeling of the wavelengths contributing to each pixel is required for precise flat fielding; (4) tilted spectra with respect to the CCD grid (the tilt varies over the field); (5) each source is dispersed into multiple orders resulting in much overlap - deep images become confusion limited; (6) zeroth-order images of compact sources can easily mimic the appearance of sharp emission features; and (7) the low resolution ($R \approx 90$ for point sources) means that only high Equivalent Width (EW) features can be discerned, while most familiar features are blends. The Space Telescope – European Coordinating Facility has done an excellent job of addressing these peculiarities with the software package *aXe* (Pirzkal et al. 2001). Armed with it and a broad-band detection image, users can extract 1D and 2D spectra that are sky-subtracted, wavelength-calibrated, flat fielded, and flux calibrated with minimum effort. Here I describe complimentary techniques I have developed to analyze WFC grism images. Specifically, I describe tools geared to finding emission-line sources and supernovae (SNe). Here I concentrate on my work with the ACS GTO team to search the Hubble Deep Field North (HDFN) for Emission Line Galaxies (ELGs) and work with the PEARS team to find SNe.

2. Initial Processing

aXe is designed so that it can work with a stack of individual dithered exposures (the FLT or CRJ images), where the grism images have not been flat fielded nor geometrically corrected. However, both flat fielding and drizzling can be very useful. Application of the F814W flat corrects most gross blemishes and removes at least half the amplitude of large-scale sky variations (after geometric correction). Spurious dark spots may remain at the blue end of some spectra, but their amplitude will be diluted if there are numerous small dithers. Their presence will have little impact on emission-line searches, while their sharpness means they are unlikely to be confused with real absorption features.

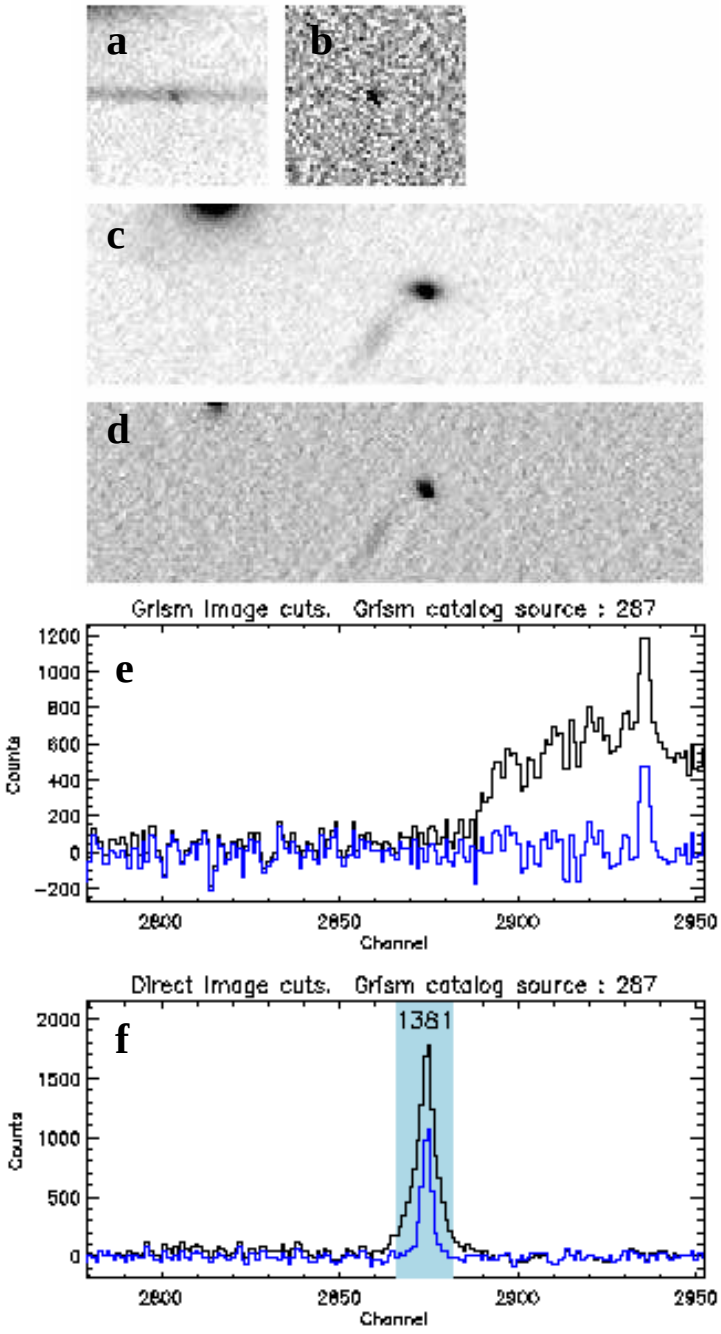

Figure 1: Steps in processing the grism and broad band detection images for finding ELGs using method B. Panels a and b show a 50×50 cutout of the grism image before and after subtracting a 13×3 median filtered version of the image. Panels c and d show cutouts of the detection image before and after the same filtering. The width of the cutout covers the full x range over which the counterpart to the source seen in panel b may reside. Panel e shows the collapsed 1D spectra of five rows centered on the emission line extracted from the grism image before (black (upper) line) and after (blue (lower) line) median filtering. Panel f is the same but for the 1D cuts through the detection image. The shaded region with identification is derived from collapsing the the *SE* segmentation image of the detection image.

Drizzle combining multiple dithered exposures is feasible as long as the dither offsets are all within 6″; then the alignment across the spectra will all be correct to within 0.5 pixels. The resultant geometrically corrected images have first order spectra that are nearly horizontal across the image, and greatly decreased spatial variation in the wavelength solution. Drizzle combining also allows improved CR rejection, especially when done with the ACS GTO pipeline *Apsis* (Blakeslee et al. 2002).

A mask is used to mark or remove the zeroth order images. First the zeroth order sources in the grism image are matched to those in a broad band detection image. The sources are found with *SExtractor* (hereafter *SE*; Bertin & Arnouts, 1996) which is used to catalog the sources in both the detection and grism images. Only compact sources are matched. Their positions are used to define a linear transformation between the detection image and the zeroth order. The scaling ratio between the matched detection and zeroth order images is determined and used to model which pixels to mask. In the HDFN the images in F775W and F850LP are typically 32 and 21 times brighter, respectively, in count rate than the zeroth order counterparts. This scaling ratio is used to determine which pixels will have zeroth order counterparts that are brighter than sky noise level. The position of these pixels in the detection image are transformed to populate a mask for the grism image which is then grown by three pixels to account for the slight dispersion in the zeroth order. Masked pixels are set to zero at the appropriate stage of the analysis.

3. Finding Emission Line Galaxies

The ACS Science team observations centered on the HDFN consist of 3 orbits with G800L and F850LP (z_{850}) and two orbits with F775W (i_{775}). Two complimentary techniques for finding ELGs were employed on this field.

A: Search 1D spectra. *aXe* is used to extract spectra of all *SE* cataloged sources in the detection image down to $i_{775} = 26.5$ AB mag. The flux calibrated spectra are then filtered by subtracting 13 pixel median smoothed spectra leaving only sharp features. Sources with peaks having $S/N > 4$ are flagged as likely ELGs. The flagged spectra are classified by eye - broad absorption line sources are also flagged by this algorithm. These are usually M or K stars, but also include the two SNe in this field (Blakeslee et al. 2003) . The true ELGs have their emission lines fitted with Gaussians to derive line wavelength and flux.

B: Search 2D grism image. The basis of this method is the observation that most emission line sources appear to correspond to compact knots, not necessarily at the center of galaxies. Here we find the line emission in the grism image first and then pinpoint the emitting sources in the detection image, as illustrated in Fig. 1. Sharpened versions of both the grism and direct images are made by subtracting a 13×3 median smoothed version of themselves. In the grism image, this effectively subtracts the continuum and removes cross-dispersion structure. This image is then cataloged with *SE*. Ribbons, typically covering five rows, centered on the y position of each source are extracted from both the sharpened grism and direct images. Since the dispersed spectrum lies primarily to the right of the direct image, the extracted ribbons extend more to the left so that all possible sources that could have created the emission line are in the direct ribbon. The ribbons are collapsed down to 1D spectra and cross correlated after the regions beyond ± 13 pixels from the source in the grism image are set to 0.0. This is done so they do not contribute to the cross-correlation amplitude. Any knot within the detection ribbon will produce a peak in the cross-correlation spectrum. The position of the peak yields the offset between the knot and the line emission in the grism image. Using the wavelength solution for the grism, in principle one could derive the line wavelength from this offset. Instead, final measurements of the emission line quantities are obtained from 1D spectra of each knot extracted with *aXe* using the cross-correlation determined position of the star forming knot. As with method A, the emission line properties are measured with Gaussian fits to the spectra.

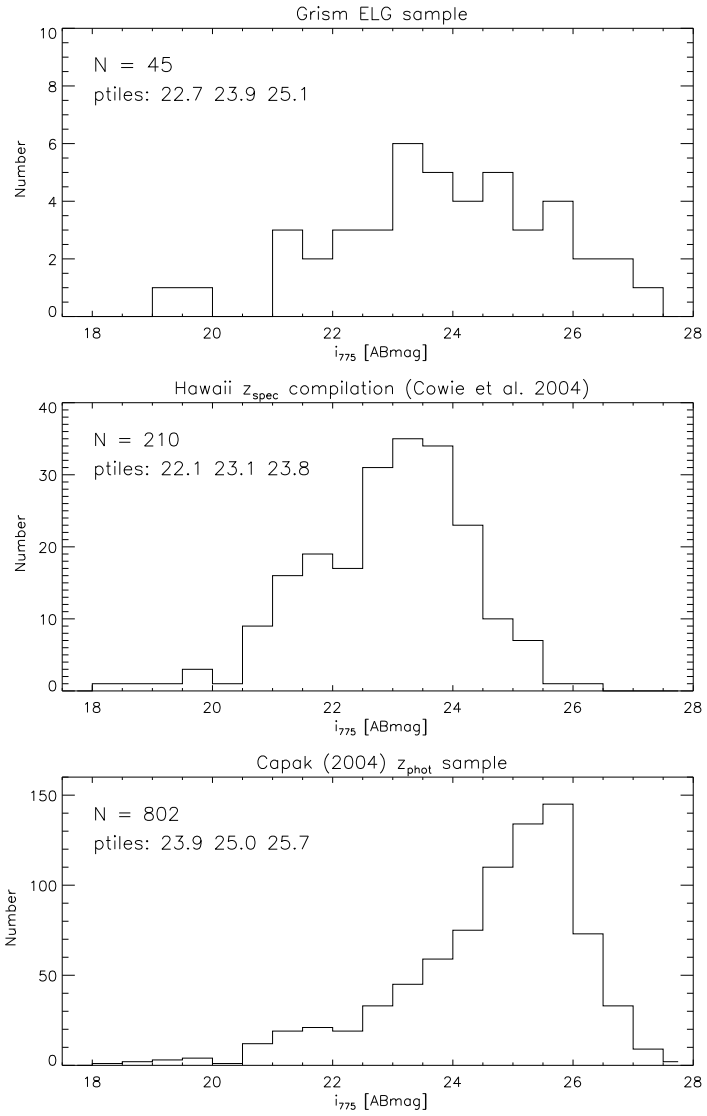

Figure 2: Histogram of i_{775} AB magnitudes of the grism selected ELG sample in the HDFN (top panel) compared with the spectroscopic redshift sample of Cowie et al. (2004; middle panel) and the photometric redshift sample of Capak (2004; bottom panel). In the upper left of each corner we report the total number of sources in the sample and the 25th, 50th (median), and 75th percentile i_{775} AB magnitudes.

In the HDFN field we found 30:39 ELGs with methods A:B. For the most part, the same galaxies are found; 7:16 ELGs were uniquely found with methods A:B. For three ELGs we identified multiple emission line knots with method B. Figure 2 compares the i_{775} apparent magnitude distribution of the merged list of ELGs from our analysis compared to galaxies with spectroscopic and photometric redshifts in the same field. The grism ELGs, found with three orbits of *HST* time, are on average fainter than the galaxies with spectroscopic redshifts gathered over several years from the Keck 10m telescopes.

4. Line Identification

Line identification is a major concern. Only seven of the ELGs in HDFN have two emission lines in our data. In those cases the lines can be identified using the ratio of wavelengths which remains invariant with redshift. However one must be careful with this technique since $\lambda_{\mathrm{H}\alpha}/\lambda_{[\mathrm{OIII}]} = 1.3138$ is close to $\lambda_{\mathrm{H}\beta}/\lambda_{[\mathrm{OII}]} = 1.3041$. A one pixel (42Å) uncertainty in both line wavelengths could result in an incorrect line identification.

The remaining sources only have one line. The dispersion is too low to split the [O II] doublet, the [O III]4959,5007Å lines are also blended, as is Hα and the [N II] doublet. With only one line, at the grism's resolution, then a good first guess redshift is crucial for line identification. Drozdovsky et al. (2005) tackle this problem, in part, by looking at the size of the host galaxies. However, size alone is not a great indicator of redshift - there is little evolution in angular size for $z > 0.2$. Our approach is to use photometric redshifts as the first guess redshift. This results in line identifications for 37 of the 39 single line ELGs.

Figure 3 compares grism redshifts with spectroscopic redshifts, in panel a, and spectroscopic versus photometric redshifts in panel b. Taking the spectroscopic redshifts as "truth" results in 1/15 : 3/19 misidentified lines with methods A:B. This is similar to the error rates from photometric redshifts, as can be discerned from Fig. 3b. The dispersion about the z_{grism} versus z_{spec} unity line, excluding the outliers is 0.016:0.009 for methods A:B. Method B is probably more accurate because it better pinpoints the location of line emission. This compares to a dispersion about the unity line in z_{phot} versus z_{spec} (after clipping outliers) of 0.073, 0.107, 0.082 for z_{phot} estimates from Capak (2004), Fernandez-Soto et al. (1999), and our own photometric redshifts respectively. Thus grism redshifts are nearly an order of magnitude more accurate than photometric redshifts.

5. Finding Supernovae

The two SNe discovered in the HDFN have broad absorption features, distinctly different from Galactic stars, and are easily visible in our grism spectra (Blakeslee et al. 2003) demonstrating the viability of grism surveys for SNe searches. The PEARS team has obtained 200 orbits of *HST* time primarily to characterize high-redshift objects in the two GOODS fields using the WFC and G800L grism. An additional aim is to search for SNe on a rapid turn-around basis. The total exposure time at each pointing/roll angle is about twice as long as the HDFN observations described above. However, only shallow broad-band images are obtained concurrently with the grism exposures. These are used to align the grism images to the astrometric grid of the GOODS fields. But they are not as deep as the grism image, hence they may not reveal SNe. So although *aXe* spectra are generated of the prior GOODS cataloged sources, they are not useful for finding SNe at later epochs.

What is needed is a method to find SNe using only the grism images. To this aim I have developed an IDL package *SHUNT* (Supernovae Hunt) to find and classify the first order spectra of all compact sources in a grism field. As before, the starting point is geometrically corrected, combined grism images. Since most source cataloging codes (i.e. *SE*) have been developed to find compact blobs, they do not work so well for finding grism spectra which are very extended, often at near the noise level of the image. Rather than optimizing

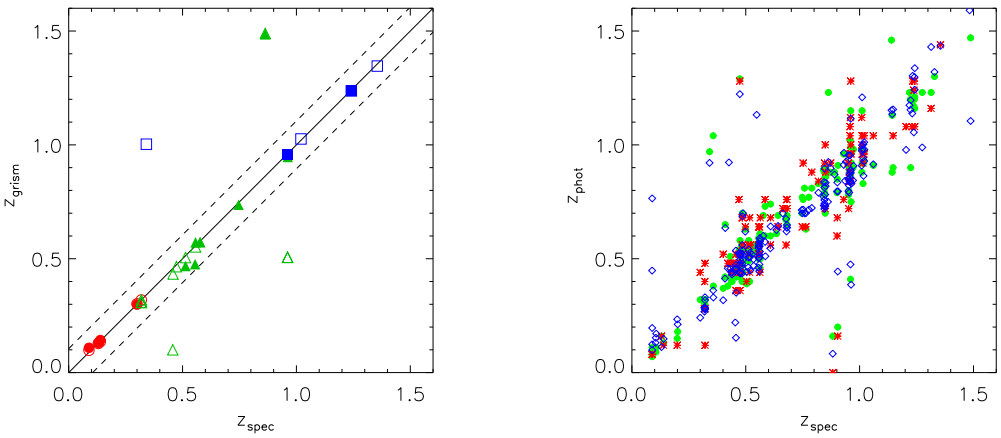

Figure 3: Comparison of grism redshifts (left) and photometric redshifts (right) with spectroscopic redshifts from Cowie et al. (2004). In the left panel, the unity relationship is shown as a solid line, sources outside the dashed lines at $\Delta z = \pm 0.105$ are outliers. Only photometric redshifts were used for the first guess redshift. If spectroscopic redshifts are used as priors there is still one outlier. Here, measurements from method A are shown with solid symbols, measurements from method B are shown as open symbols. The symbol shape and color indicate the grism line identification: Hα emitters are (red) circles, [O III] emitters are (green) triangles and [O II] emitters are (blue) squares. In the right panel the photometric redshifts from Cowie et al. (2004), Fernandez-Soto et al. (2004), and our measurements are shown as (green) filled circles, (red) asterisks, and (blue) hollow diamonds respectively.

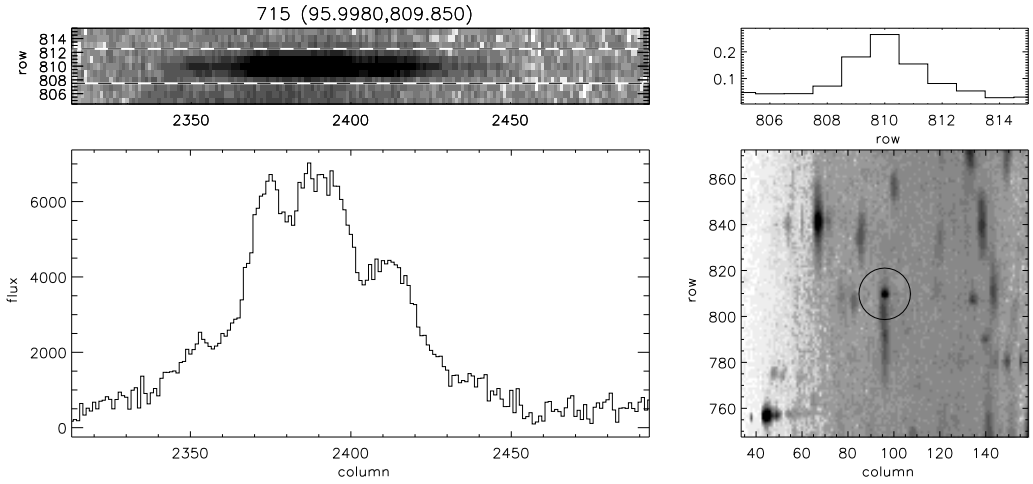

Figure 4: Example of a supernova identified with *SHUNT*. The top left panel shows the geometrically corrected grism image. The bottom left panel shows the extracted 1D spectrum found by collapsing the above 2D image between the dashed lines. The top right panel shows a 1D cut along the cross dispersion of the spectrum. The bottom right panel shows the squashed grism image with the source identified.

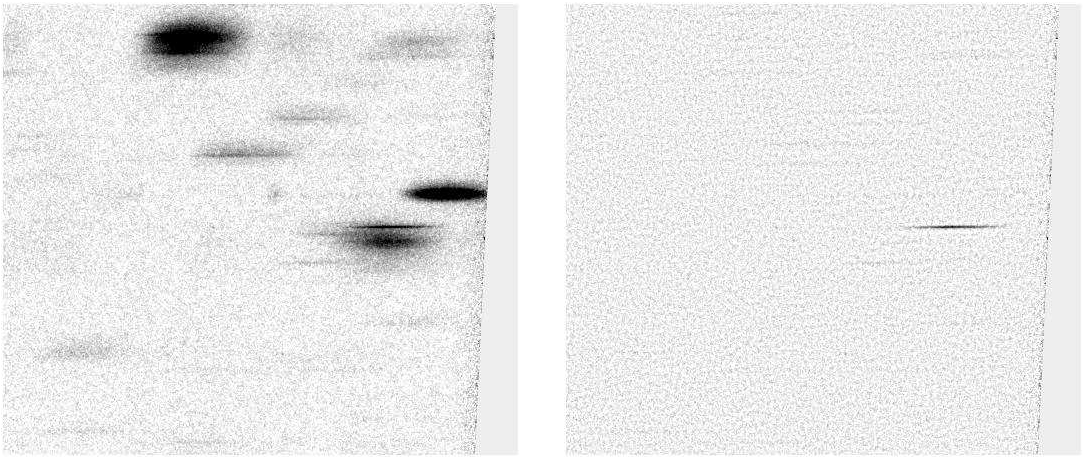

Figure 5: Portion of a grism image containing a faint late type SN. Panel a (left) shows the geometrically corrected and combined image. Panel b (right) shows the same image after subtracting the *aXe* model image based on prior GOODS observations, and then subtracting a smoothed version of the residuals. The SN now clearly stands out.

the code to fit the data, *SHUNT* makes the data suit the code by squashing (rebinning) the image 25×1 per output pixel before cataloging with *SE*. This results in first order spectra being close to critically sampled in the x direction. The resultant catalog is filtered to remove small sources (typically zeroth order images) and extended sources (galaxies). The remaining ~ 250 sources are then classified. Five rows centered on each source are collapsed to form a spectrum (which is not wavelength calibrated). Classification is by eye where the classifier (me) examines figures such as Fig 4 showing the 2D spectrum, the collapsed 1D spectrum, a cross dispersion trace and the squashed grism image. Each source is classified as either SN, unidentified absorption spectrum, probable M or K star, break spectrum, emission line source, featureless, non-first order spectrum, or spurious (the order is of decreasing interest, and roughly of increasing occurrence rate). Direct postage stamp images from GOODS (or the shallow broad-band images) are generated with a rectangular error box plotted which should contain the source. An empty error box in the prior GOODS image is a second indication of a transitory source. It typically takes 0.5 to 1 hour to classify all the objects in a field.

One problem with this approach is that it can miss SNe blended with galaxy spectra. This is more likely to occur for late time SNe spectra which can have low S/N and/or be featureless at grism resolution. An example is shown in Fig 5. One way such objects can be found is to subtract model spectra of the sources cataloged by GOODS using *aXe* v1.5 (Kümmel, this volume). The models are very good but not perfect. However, subtraction of the smoothed residuals is sufficient to isolate faint transient object spectra from the model residuals.

6. Summary

The ACS grism produces amazingly rich datasets. While the data are somewhat difficult to interpret, tools have been developed to make the most use of these data. Public access tools like *aXe* are readily available to remove most of these complications and extract 1D and 2D spectra. Here I have shown how some common manipulations of the data (such as geometric correction and flat fielding) allow interesting sources such as emission line galaxies and supernovae to be efficiently found.

Acknowledgments. Many members of the ACS and PEARS Science teams as well as others have contributed to the work presented here. I particularly thank Zlatan Tsvetanov, Holland Ford, Caryl Gronwall, John Blakeslee, Peter Capak, Sangeeta Malhotra, Norbert Pirzkal, Chun Xu, Txitxo Benitez, James Rhoads, Jeremy Walsh, and Martin Kümmel.

References

Blakeslee, J. P., Anderson, K. R., Meurer, G. R., Benitez, N., & Magee, D. 2002, ASP Conf. Ser. 295: ADASS XII, p. 257

Blakeslee, J. P., et al. 2003, ApJ, 589, 693

Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393

Capak, P. L. 2004, Ph. D. Thesis, U. Hawaii

Cowie, L. L., Barger, A. J., Hu, E. M., Capak, P., & Songaila, A. 2004, AJ, 127, 3137

Drozdovsky, I., Yan, L, Chen, H.-W., Stern, D., Kennicutt, R., Spinrad, H., & Dawson, S. 2005, AJ, 130, 1324

Fernández-Soto, A., Lanzetta, K. M., & Yahil, A. 1999, ApJ, 513, 34

Pirzkal, N., Pasquali, A., & Demleitner, M. 2001, ST-ECF Newsletter, 29, "Extracting ACS Slitless Spectra with aXe", p. 5 (`http://www.stecf.org/instruments/acs`)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Calibration of ACS Prism Slitless Spectroscopy Modes

S. S. Larsen, M. Kümmel and J. R. Walsh

ST-ECF, Karl-Schwarzschild-Str. 2, D-85748 Garching bei München, Germany

Abstract. The Advanced Camera for Surveys is equipped with three prisms in the Solar Blind (SBC) and High Resolution (HRC) Channels, which together cover the 1150 – 3500 Å range, albeit at highly non-uniform spectral resolution. We present new wavelength- and flux calibrations of the SBC (PR110L and PR130L) and HRC (PR200L) prisms, based on calibration observations obtained in Cycle 13. The calibration products are available to users via the ST-ECF/aXe web pages, and can be used directly with the aXe package. We discuss our calibration strategy and some caveats specific to slitless prism spectroscopy.

1. Introduction

The Advanced Camera for Surveys is currently the only instrument on *HST* which provides spectroscopy in the UV-optical range. With STIS being unavailable (at least for the moment), the interest in the ACS spectroscopic modes has increased substantially. An extensive calibration effort was undertaken in Cycle 13 in order to provide much improved wavelength- and flux calibrations for the prism modes which had seen little use in earlier cycles. The ACS has two prisms (PR110L and PR130L) installed in the Solar Blind Channel (SBC), both covering the wavelength range from roughly 1200Å – 2000Å. The main difference between the two SBC prisms is that the sensitivity of PR110L extends below the geocoronal Ly α, while that of the PR130L does not. The sensitivity below 1216 Å results in a significantly higher sky background for the PR110L. The one prism (PR200L) in the High Resolution Channel covers the range $\sim$ 1800Å– 3500Å. In this paper we present the calibration observations and discuss the derived trace-, wavelength and flux calibrations. The calibrations presented here are implemented in configuration files for the aXe package (Kümmel et al., these proceedings) and are made available to users via the aXe web pages (`http://www.stecf.org/software/axe`).

Spectroscopic observations with the ACS prisms share many similarities with the G800L grism in the WFC and HRC channels, but there are also some differences. In both cases the object spectrum (or spectra) may fall anywhere on the detector, and the reference point for the wavelength scale is typically established using a direct image taken immediately before (or after) the prism exposure. The direct image is even more crucial for the prism modes, which do not show a zeroth order spectrum (nor any higher- or negative orders). Contrary to the case of grism spectroscopy, the wavelength scale of the prism spectra is highly non-linear, with the spectral resolution decreasing towards longer wavelengths (see below). For the PR200L, this causes a "red pile-up" with wavelengths between 4000 Å and 10000 Å being compressed into only 7 pixels. This pile-up can be particularly troublesome for red objects, where light from the diffraction spikes and outer halo of the PSF may contaminate the bluer parts of the spectrum. These effects are still poorly quantified. No such pile-up is seen for the SBC prisms, due to the lack of sensitivity at wavelengths $>$ 2000Å for the MAMA detector.

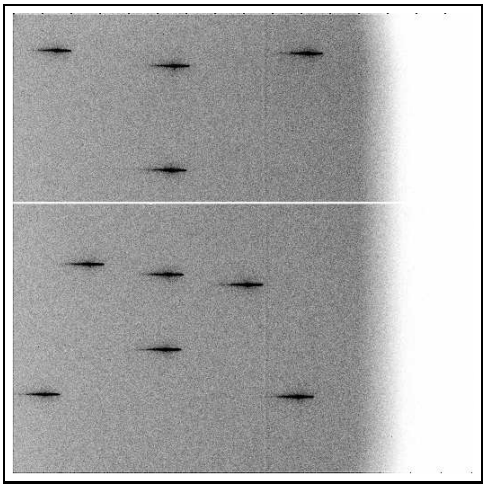

Figure 1: Sum of the 10 pointings obtained for PG1322+659 for the PR110L. Note the vignetting at the right-hand side of the field. A few dead rows at $Y \approx 600$ are seen as a horizontal white line.

2. Observations

The relatively low spectral resolution of the prisms, particularly at longer wavelengths, and the high spatial resolution of the SBC and HRC cameras limit the choice of suitable wavelength calibrators. The calibration targets must be compact *and* have strong emission features in the ultraviolet spectral region. For the SBC, an additional constraint was the "bright object protection" limit. We eventually settled on a combination of a planetary nebula in the Large Magellanic Cloud (LMC-SMP-79) and a QSO at redshift $z = 0.836$ (Q0037-3544) for the PR200L, and two QSOs at redshifts $z = 0.168$ (PG1322+659) and $z = 0.098$ (PG1404+226) for the SBC prisms. The PN was known from existing STIS spectra to have strong C III] and [C II] emission lines and be relatively compact (though not point-like). The QSOs were selected from the compilation of Véron-Cetty & Véron (2003), which lists 48921 quasars and makes it possible to select a QSO with emission features (Ly α, C IV etc.) at essentially any desired wavelength. We required the calibration targets to have existing *HST* spectra (either from FOS or STIS) and low reddenings ($A_B < 0.1$ mag), but even with these additional constraints it was not difficult to find suitable QSO targets. For the flux calibration we observed two white dwarf standards, WD1657+343 and LDS749B. STIS spectra of these stars were kindly provided by R. Bohlin.

Each target was observed at several positions across the detector in order to map spatial variations in the trace- and wavelength solutions, as well as any large-scale sensitivity variations. For each prism, one wavelength calibration target was observed at 9 or 10 positions (taking 2 orbits) and the other at 5 positions (one orbit). The flux standards were observed at 5–6 different positions. Each prism exposure was preceded by a direct image through the F165LP filter (for the SBC) or F330W (for the HRC). The F165LP filter was chosen in order to avoid conflict with the bright object protection limit for the QSO targets. However, for one of the flux standards we obtained direct images in both the F122M and F165LP filters in order to check for any offsets between exposures in the two bands. Such offsets were indeed found (see below).

The coverage of the field is illustrated in Figure 1, which shows a sum of the 10 PR110L exposures obtained for PG1322+659. The MAMA detector has a few dead rows at $Y \approx 600$ which show up as a horizontal white line in the figure. To avoid these, the default aperture center is at (512, 400). Note that there is significant vignetting at the right-hand side of the

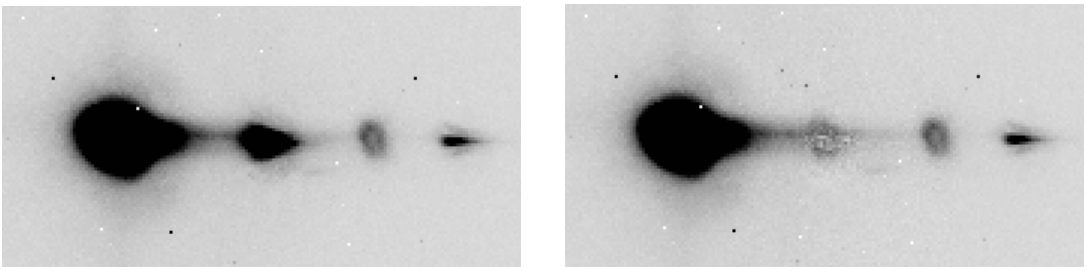

Figure 2: PR200L image of LMC-SMP-79. Left: raw image, Right: after subtraction of contaminating object. Note that the nebula is resolved at the HRC resolution. From the left to the right, the four main "features" in the spectrum are: the red pile-up, [C II] 2325Å/ contaminating star, C III] 1909Å, and another contaminating star.

field, affecting ∼200 image columns. A similar phenomenon is seen for the HRC/PR200L, but affecting the left-hand side of the field.

Two difficulties were encountered with the data for LMC-SMP-79. First, with a diameter of about 0.″3, the PN was slightly resolved at the resolution of the HRC. Second, by an unfortunate coincidence the prism spectra of the nebula were contaminated by two stars in the field, one of which affected the [C II] 2325 Å line (Figure 2). Fortunately, we found that the contaminating object could be removed by scaling and subtracting the spectrum of a star of similar spectral type located elsewhere in the field. This procedure worked well enough that we could recover the [C II] 2325 Å line. Figure 2 also illustrates the PR200L "red pile-up", which is the dominant feature at the left (red) end of the spectrum.

3. Wavelength- and trace solutions

The spectral trace- and wavelength solutions are defined with respect to a reference position, $(X_{\rm ref}, Y_{\rm ref})$, which was measured by running the SExtractor code (Bertin & Arnouts 1996) on the direct images. The trace descriptions were then derived by measuring the centroid along the image columns (corresponding to the spatial direction) for the flux standard spectra, and fitting a straight line (for the PR110L and PR200L) or 2$^{\rm nd}$-order polynomial (for PR130L) to the offset $\Delta Y = Y_{\rm trace} - Y_{\rm ref}$ vs. $\Delta X = X_{\rm trace} - X_{\rm ref}$. In all cases the dispersion direction of the prism spectra was found to be aligned with the image rows to within 1 degree, towards positive ΔX for PR200L and negative ΔX for the SBC prisms. The variation in the offsets and slopes $\Delta Y/\Delta X$ of the prism spectra was fitted with 2-dimensional 1st order polynomials as a function of position on the detector (not reproduced here due to space constraints).

Wavelength solutions were assumed to be of the form used by Bohlin et al. (2000), i.e.

$$\lambda = a_1 + \frac{a_2}{(\Delta X - a_0)} + \frac{a_3}{(\Delta X - a_0)^2} + \frac{a_4}{(\Delta X - a_0)^3} + \frac{a_5}{(\Delta X - a_0)^4} \qquad (1)$$

where each of the coefficients $a_0 \ldots a_5$ can be a function of $(X_{\rm ref}, Y_{\rm ref})$. In order to define the wavelength solutions, the prism spectra were first extracted with aXe using a configuration file where the "wavelength" scale was simply the pixel offset, ΔX, along the trace. The aXe spectra were then converted to IRAF format, and the ΔX values for the various emission lines were measured with the SPLOT task, using a Gaussian fit. We then solved directly for the full, spatially dependent, wavelength solution. In the PN spectrum we could measure the C III] 1909Å and [C II] 2325Å lines and in the spectrum of Q0037-3544 we could measure Lyα, C IV and C III] redshifted to 2233Å, 2844Å and 3503Å, thus providing 5 wavelength sampling points for the PR200L. For the SBC prisms, we detected Ly α redshifted to 1420Å

Table 1: Wavelength solution terms

SBC:	PR110L	PR130L
a_0	$-95.70 - 0.00662 \times X_{\rm ref} + 0.01760 \times Y_{\rm ref}$	$-80.61 - 0.00711 \times X_{\rm ref} + 0.01547 \times Y_{\rm ref}$
a_1	$1049.39 + 0.02105 \times X_{\rm ref} - 0.01867 \times Y_{\rm ref}$	$1072.71 + 0.01249 \times X_{\rm ref} - 0.01699 \times Y_{\rm ref}$
a_2	2.11506×10^4	-1.49676×10^4
a_3	9.23749×10^6	1.94800×10^6
a_4	4.93403×10^8	5.48481×10^7
a_5	1.24585×10^{10}	1.11636×10^9

HRC:	PR200L
a_0	$145.755 + 0.00872 \times X_{\rm ref} - 0.01925 \times Y_{\rm ref}$
a_1	$1147.24 - 0.03042 \times X_{\rm ref} - 0.00782 \times Y_{\rm ref}$
a_2	8.11002×10^4
a_3	-8.28929×10^5
a_4	-4.65350×10^6
a_5	1.68888×10^8

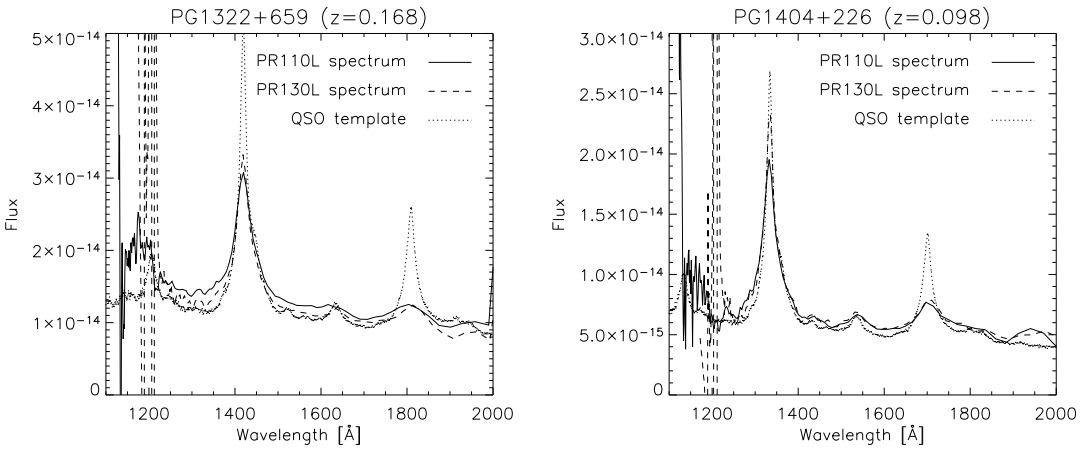

Figure 3: SBC PR110L and PR130L spectra of PG1322+659 (left) and PG1404+226 (right). A scaled template QSO spectrum (from Zheng et al. 1997) is shown for comparison.

and 1335Å in the spectra of PG1322+659 and PG1404+226 respectively, as well as C IV redshifted to 1701Å for PG1404+226. The spectra of PG1322+659 had too low resolution at the wavelength of C IV (1806Å) to allow useful measurements of this feature. Thus, we have 3 wavelength sampling points for the calibration of the SBC prisms. The resulting wavelength solutions are reproduced in Table 1. Note that we have adopted the higher order terms from Bohlin et al. (2000), i.e. a_3, a_4 and a_5 for the PR200L, and $a_2 \ldots a_5$ for PR110L and PR130L. However, the a_2 term we derive for the PR200L is quite similar to that derived by Bohlin et al. (81100 vs. 83000).

The trace- and wavelength solutions derived here assume that direct images are taken in the F165LP filter (for the SBC) or in F330W (for the HRC). For the SBC, images taken in different filters show offsets of up to 12 pixels, so it is important to correct any coordinates measured on images taken in filters other than F165LP for these offsets. A preliminary analysis of these offsets, using our data and data from program 9563 (P.I.: G. De Marchi), suggests that the following corrections should be applied (all in pixels): F115LP: $(\Delta x, \Delta y) = (-3.74, 12.13)$; F122M: $(\Delta x, \Delta y) = (-3.74, 12.07)$; F125LP: $(\Delta x, \Delta y) = (0.97, -2.44)$; F140LP: $(\Delta x, \Delta y) = (2.88, -8.48)$ and F150LP: $(\Delta x, \Delta y) = (-0.31, 1.26)$.

Sensitivity functions were derived for each prism by dividing the prism spectra of the flux standards with the standard STIS spectra, and fitting the ratio with a spline function.

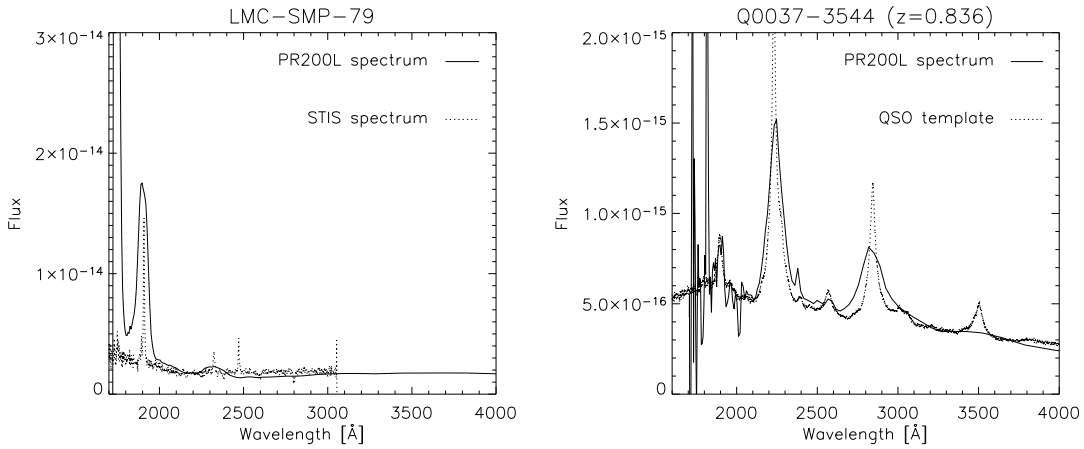

Figure 4: HRC PR200L spectra of LMC-SMP-79 (left) and the quasar Q0037-3544 (right). For comparison, a STIS spectrum of LMC-SMP-79 (courtesy L. Stanghellini) and a scaled quasar template spectrum are shown.

The spectra of the flux standards were extracted using an extraction box height of $\pm 0\farcs5$. A flat field cube (Kümmel et al., these proc.) has been defined for the SBC prisms, although not including any wavelength-dependency. This shows sensitivity variations across the SBC detector amounting to a few percent.

4. Consistency checks

Figures 3 and 4 show the SBC and HRC prism spectra of the wavelength calibration targets, extracted with aXe, compared with a template QSO spectrum (Zheng et al. 1997) and the STIS spectrum of LMC-SMP-79 (courtesy of L. Stanghellini). The QSO template spectrum has been adjusted to the redshift of the corresponding targets and scaled to match the data. No reddening correction has been applied, but the observed QSOs all have $A_B < 0.1$ mag. The prism observations appear in good agreement with the reference spectra, both as far as the wavelength scales and flux calibrations are concerned. While comparison of aXe-extracted prism spectra of the *flux* standards with the STIS reference spectra shows agreement to about 5%, this agreement may be somewhat deceptive, since the sensitivity files were of course derived from those observations in the first place. The comparison in Figures 3 and 4 thus provides a welcome independent check, for objects which have rather different spectral energy distributions than the WD standards. While the absolute scalings of the reference spectra here are arbitrary, the overall shape of the prism spectra appears correct, even at longer wavelengths where the wavelength solutions become increasingly uncertain. However, note that the decrease in spectral resolution at longer wavelengths makes it difficult to detect even relatively strong features, such as those seen in the QSO spectra.

The wavelength scale is generally accurate to better than a pixel. The dispersion of the SBC prisms is about 2 Å pixel^{-1} at 1200 Å, but degrades to about 10 Å pixel^{-1} at 1600 Å and 30 Å pixel^{-1} at 2000 Å. The PR200L has a dispersion of about 5 Å pixel^{-1} at 1800 Å and nearly 40 Å pixel^{-1} at 3000 Å. So while the three prisms together cover the full wavelength range from 1150 Å– 3500 Å, the spectral resolution is highly non-uniform.

Using the sensitivity curves derived here, we calculate the total system throughput for each prism, as shown in Figure 5. The PR200L reaches a peak efficiency of about 10%, while the SBC prisms peak at 3–4%. These estimates are simply the measured count rates vs. the

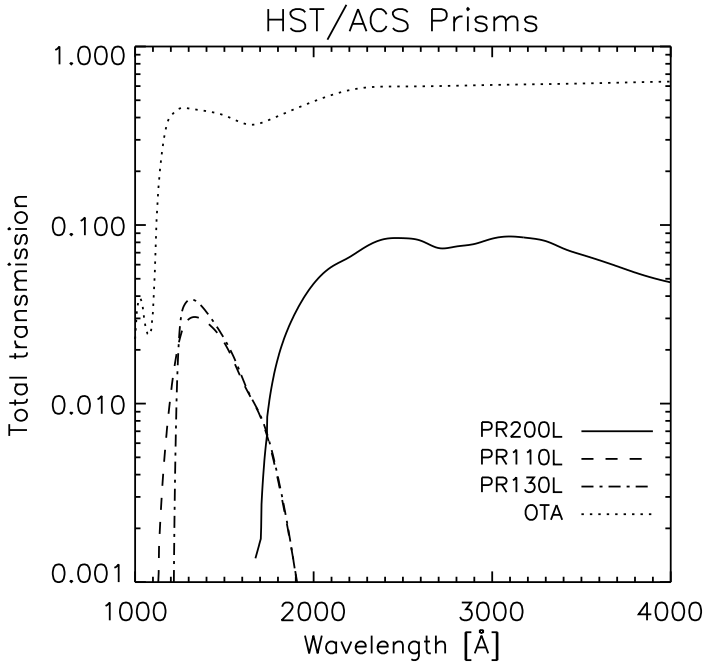

Figure 5: Total transmission of the ACS prisms relative to the "perfect telescope", based on the sensitivity curves derived here. The dotted line indicates the throughput of the OTA.

maximum possible rates, assuming a telescope aperture of 4.52 m^2. The throughput curves thus include all the effects of mirror coatings, detector sensitivity etc. For comparison, the throughput of the Optical Telescope Assembly (OTA) alone is also shown.

5. Summary

The SBC and HRC prisms are now fully calibrated, both in wavelength and flux. The calibration products are available for use with the aXe package, and can be downloaded from the aXe web pages. Future calibration observations will aim to monitor the stability of the prism modes and possibly provide various improvements, such as quantifying the effect of scattered light from the red pile-up in the PR200L. A better understanding of the general properties of the SBC and HRC detectors, such as CTE effects, may also lead to improvements in prism spectroscopy.

References

Bertin, E., & Arnouts, S., 1996, A&A Suppl., 117, 393

Bohlin, R. C., Hartig, G., & Boffi, F. R., 2000, *Instrument Science Report* ACS 00-01 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Véron-Cetty, M. P., & Véron, P., 2003, A&A, 412, 399

Zheng, W., Kriss, G. A., Telger, R. C., Grimes, J. P., & Davidsen, A. F., 1997, ApJ, 475, 469

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

WFPC2 Status and Calibration

J. Biretta

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. We briefly summarize the status of WFPC2 as of late 2005, and review calibration results from the last few years.

1. Introduction

The Wide-Field / Planetary Camera II was built at the Jet Propulsion Laboratory in Pasadena, CA under the leadership of John Trauger. It was installed into HST in December 1993 and is the oldest imager on the spacecraft. With over 50,000 science images to its credit, it accounts for roughly half the science images taken by HST. Even after the installation of more modern cameras, it continues to represent about 10% to 15% of the HST science program. Its capabilities for wide-field UV and narrow band imaging remain unique, as does its availability for long-term astrometric and monitoring projects. Due to its wide field of view it also makes significant contributions as a parallel imaging instrument.

After 12 years on-orbit it continues to operate relatively well. A new electronic anomaly was recently discovered in the WF4 channel, and we discuss initial assessments of its impact and possible mitigation strategies. Long-term radiation damage causes a number of effects, including increasing numbers of hot pixels and elevated dark current, but these are not serious issues. Charge Transfer Efficiency (CTE) issues in the CCDs, also caused by radiation damage, continue to slowly increase and represent perhaps the most significant lien against the instrument's performance. We briefly summarize results of calibrations carried out over the last few years in the areas of CTE, photometry, and geometric distortion. Finally we discuss potential "Close Out" projects which might further refine calibrations and increase the archival value of WFPC2 data. Herein we highlight only a few recent developments; for a more complete picture of the WFPC2 calibration the reader is referred to the WFPC2 Instrument Handbook, the HST Data Handbook, numerous Instrument Science Reports, and excellent reviews in prior Calibration Workshops (e.g. Whitmore 1997, Koekemoer 2002, Whitmore 2002).

2. New Anomaly in the WF4 CCD

A new anomaly was recently discovered in the electronics of the WFPC2 WF4 channel. It is characterized by sporadic images with either low or zero bias level from the WF4 CCD. Since the calibration pipeline automatically subtracts the bias level, images with low bias levels are not immediately obvious, however more careful examination will reveal faint horizontal (X-direction) stripes in the background and low target brightnesses. Images with zero bias level are much more obvious – these images are mostly blank, although bright pixels such as cosmic rays, strong hot pixels, and very bright targets are sometimes visible. So far the other three camera channels appear normal and completely unaffected by this anomaly.

The anomaly was first identified in routine dark calibration frames taken on 23 September 2005, but it is clear that the anomaly has been slowly developing for several years. Figure 1 shows the bias levels of WFPC2 gain=7 images taken between 2003 and late 2005.

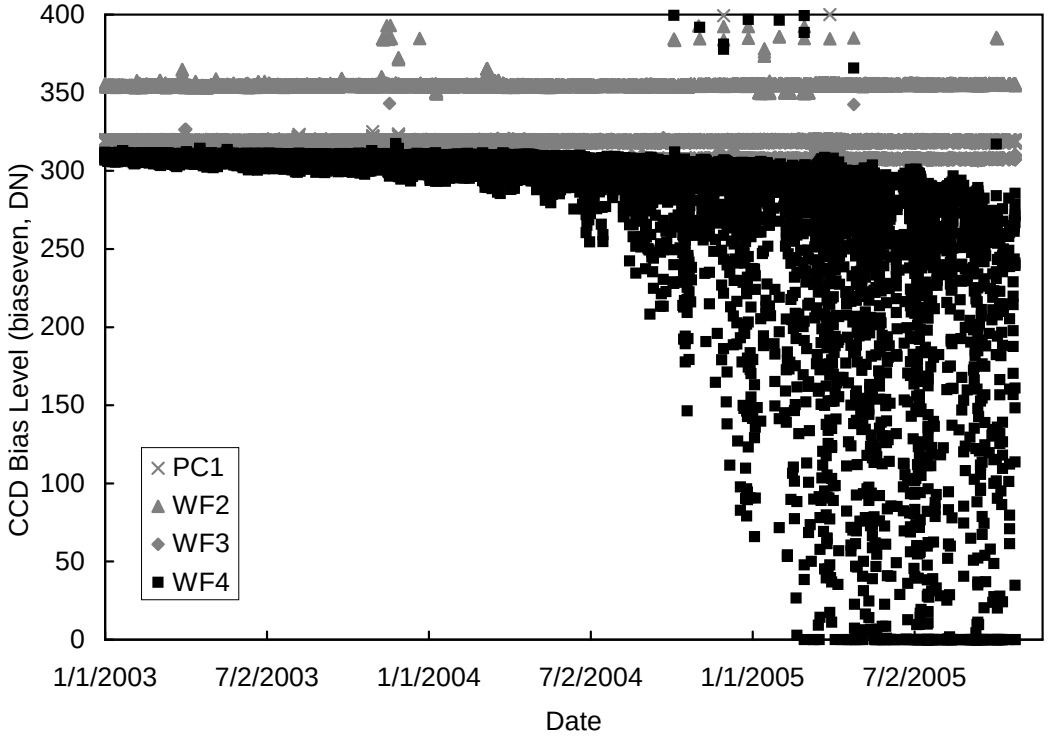

Figure 1: Bias levels of WFPC2 images between 2003 and the current epoch at A-to-D converter gain 7. Each point represents a single WFPC2 image. Bias levels in PC1, WF2, and WF3 are highly stable, whereas WF4 becomes becomes increasingly unstable and tends toward low values. (Sporadic high points, mostly in WF2, are highly saturated images.)

As can be seen, sporadic images with bias levels significantly below the normal level ∼311 DN begin during 2004, and by February 2005 the first zero bias images are seen. By the end of 2005 all images in WF4 have below-normal bias levels, and ∼20% have zero bias level and are blank. The earliest traces of the anomaly appear in early 2002, near the time of Service Mission 3B, when WF4 first shows sporadic images with bias levels a few DN below normal.

Images with low bias level suffer two primary issues. Firstly, the images contain faint horizontal streaks with an amplitude of about 1 DN RMS. Fortunately, these can usually be removed with appropriate spatial filtering of the image. Secondly, the photometry is corrupted in the sense that targets appear too faint. The strongest photometric effect is seen for faint targets at low bias levels. At gain=7 and for bias levels approaching zero, faint targets are about 40% too faint. At gain=15 targets can be as much as 70% too faint. This photometric anomaly appears to be caused by non-linearities in the signal chain at very low voltage levels, i.e. voltage levels not experienced during normal operations. Fortunately the photometric anomaly appears to be highly deterministic in the sense that a given parameter set (i.e. gain setting of 7 or 14 electrons per DN, bias level, and target signal in DN per pixel) always produces the same error, and therefore should be correctable. Observations specifically aimed at measuring the photometric effects are being carried out (calibration program IDs 10772 and 10777). Preliminary corrections for both the background streaks and photometry are discussed by Biretta & Gonzaga (2005).

Images with zero bias level are not correctable. Most of the image pixels will fall below the "zero" level of the analog-to-digital converter, and hence will register as zero DN. A few bright pixels may still be visible (cosmic rays, very bright target pixels, etc.) if they rise above the A-to-D converter zero value.

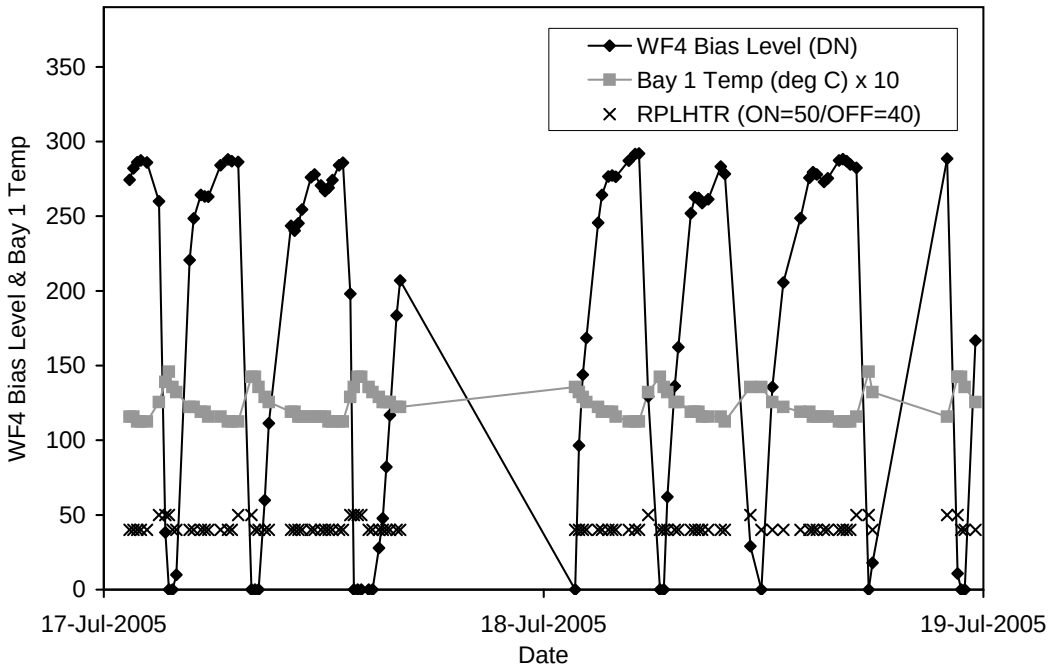

Figure 2: WF4 bias levels plotted along with Bay 1 temperature and Replacement Heater (RPLHTR) activations between July 17 and 19, 2005 to illustrate typical behavior. Note correlation of low bias levels with peaks in temperature. Temperature peaks are driven by activations of the Replacement Heaters which are indicated on the plot by RPLHTR values of 50.

The root cause of the anomaly is not yet understood at this time, but investigations are focussing on a circuit in the WF4 camera head that samples and amplifies the signal from the CCD. Scenarios involving both component failure and radiation damage induced failure seem possible.

Incidences of the anomaly appear to be strongly correlated with peaks in the temperature of the WF4 camera head electronics. The anomaly is most severe near peaks in the temperature (Figure 2), and hence a possible mitigation strategy is to reduce the temperatures and temperature swings of the electronics. The temperature swings are caused by cycling of the Replacement Heaters inside WFPC2, which serve to maintain the optical bench near 12 degrees C. The heaters are crudely controlled by comparing the Bay 1 temperature against a pair of software set points at 11 and 15 degrees C. The heaters go on when Bay 1 falls to 11 degrees C, and go off when Bay 1 reaches 15 degrees C. By reprogramming these set points, it should be possible to reduce the temperature of the WF4 camera head electronics where we suspect the problem lies. Obviously some attention must also be given to the optical bench stability. On-orbit tests of this strategy are planned for January and February 2006.

[Note added in proof: In January 2006 a test was made with the upper set point of the Replacement Heater reduced from 14.9 to 12.2 degrees C. This strategy was extremely successful. The range of WF4 bias values went from 0 to 240 DN, to about 150 to 270 DN. Most importantly there are no longer zero bias or blank images. No significant changes in the optical alignment or optical image quality were seen. Immediate plans are to leave the temperatures at these new settings, and to explore further temperature reductions.]

More detailed information on the anomaly can be found in Biretta & Gonzaga 2005, as well as later postings on the WFPC2 website. An Anomaly Review Board has been formed (chaired by Ed Cheng of GSFC) and a detailed report is expected a few months into 2006.

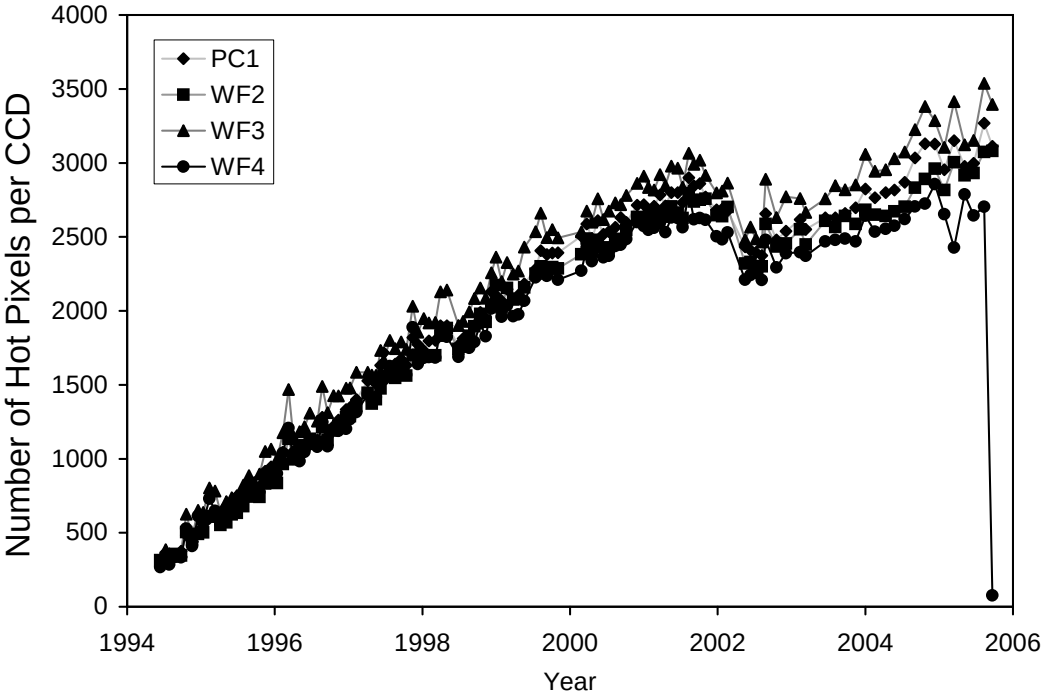

Figure 3: Number of hot pixels per CCD vs. year. Here we define a hot pixel as having a dark current in excess of 0.005 DN per second a gain 7, which is roughly the level at which a hot pixel represents a significant error in a typical 300 second science exposure.

3. Radiation Damage: Hot pixels and Dark Current

Given WFPC2's twelve years on-orbit, radiation damage is an obvious concern. Typical manifestations of damage include increased numbers of hot pixels, increased dark current, and charge transfer issues.

Figure 3 shows the number of permanent hot pixels in each CCD as a function of time. As can be seen, all the CCDs show a roughly linear growth in the number of hot pixels over time. Even at the current epoch these hot pixels represent a very small fraction of the pixels, and hence have minimal impact on science observations. Moreover, given the small fraction of hot pixels, dithering of the target position is highly effective in removing them.

Approximately once per month the CCDs are warmed to +20 degrees C both to clear the cold optical windows on the CCDs (i.e. decontaminate the windows), and to remove the hot pixels via annealing of the CCD. The plotted data in Figure 3 are from 1800 second dark frames taken immediately after each of these DECONs, and hence these represent permanent hot pixels. Near the beginning of 2002 all the CCDs show a sudden drop in the number of hot pixels. This drop appears to be associated with two events where the CCDs were left in a warm state for several days (an HST safing and SM3B). Apparently these events were able to correct additional hot pixels beyond what is typically annealed by the 6 hour DECONs.

The low-level or background dark current which is present in all pixels is illustrated in Figure 4. This shows an interesting and complex behavior. The different CCDs initially have values that are nearly a factor of three different, with WF2 being the lowest, and PC1 being the highest. All the CCDs initially show a linear increase, but around 2000 the increase is slower. Part of the complexity is caused by the fact that some of the dark current is not electronic, but rather originates in the CCD windows, perhaps as a result of scintillation. This optical component of the dark current correlates with the cosmic ray flux

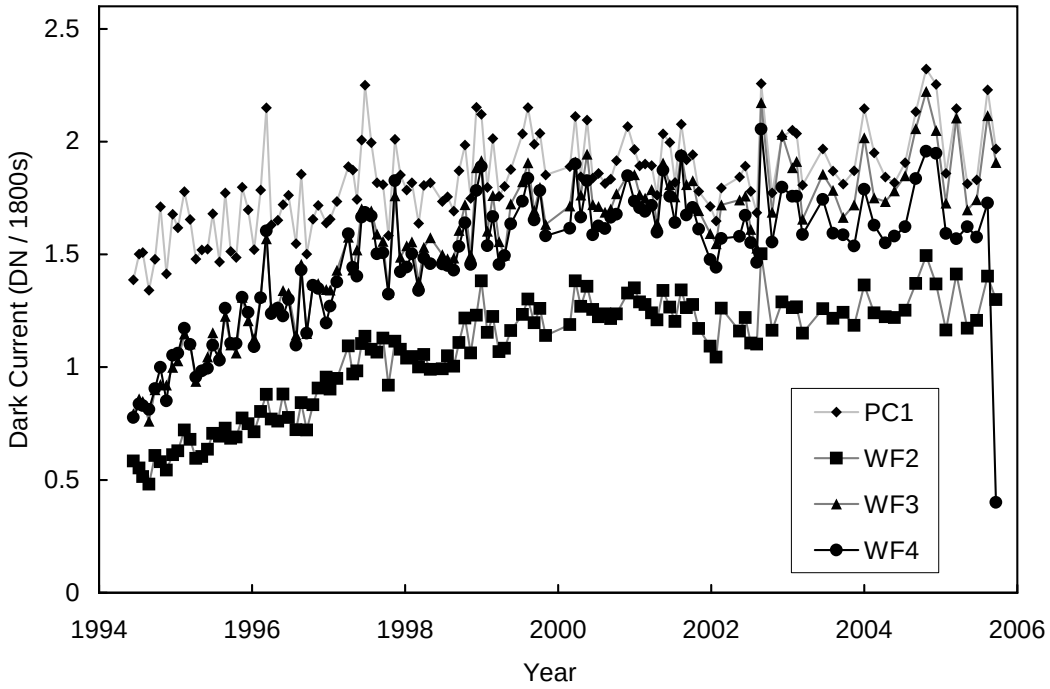

Figure 4: Low-level dark current vs year at gain 7.

in the images, and is strongest in PC1 and weakest in WF2, and probably accounts for the different dark currents in the different CCDs. Some of the apparent leveling-off after 2000 may also be related to the long warm events near the beginning of 2002. But even after these considerations, it is not apparent why the dark current increase is slower at current epochs than after the initial installation of WFPC2. One speculative cause is CTE effects, which might trap some of the dark current at late epochs. In any case, the dark current does not appear to have become a significant issue for WFPC2 image quality.

Both Figures 3 and 4 show WF4 having low data points in late 2005. This is due to the WF4 anomaly producing low counts and sometimes blank images. In fact the anomaly was first discovered while updating Figure 4.

4. Radiation Damage: Charge Transfer Issues

The most significant radiation-induced issue for WFPC2 are problems with the charge transfer efficiency (CTE) in the CCDs. The CCDs are read out by transferring the charge in each pixel across the CCD chip to a single readout amplifier; CTE problems arise when some portion of the charge is left behind during this readout process. The problem is well-known and there are many references (e.g., Whitmore, Heyer, & Casertano 1999; Dolphin 2000). CTE is especially troublesome because it disproportionately impacts faint targets, and hence will ultimately limit the ability to detect faint objects. It also increases with time as radiation damage grows. While it was once only a significant issue for faint targets on faint backgrounds, it now has a significant photometric impact on all brightness levels.

There has not been a detailed trending of its photometric progress during the last few years, but we believe the previous relations and web-tools are still appropriate (e.g., see `http://www.stsci.edu/instruments/wfpc2/wfpc2_cte/wfpc2_cte_calc.html`). For example, at the end of year 2005 we predict that a moderately bright target (2000 electrons total) in a typical 300 second F555W exposure near the CCD center will lose about 10% of its counts or require a correction of 0.10 magnitudes. Even a very bright target approaching

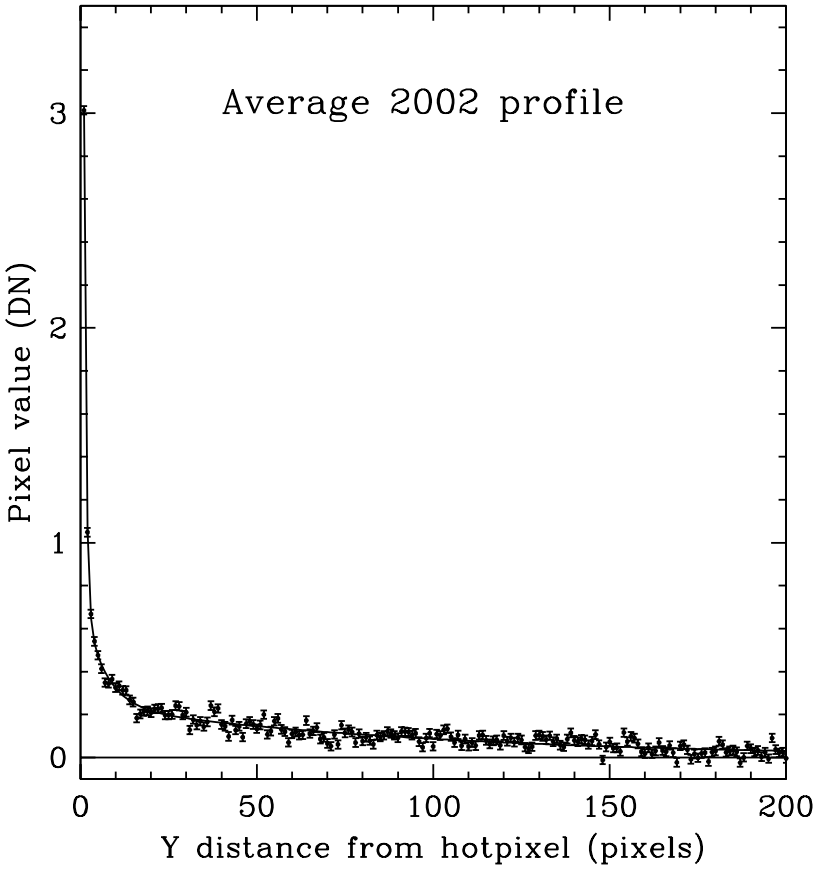

Figure 5: Average hot pixel tail at epoch 2002. Corresponding hot pixel brightness is 310 DN.

saturation (55,000 electrons total) loses about 5% of its counts, and will require correction for accurate photometry. The situation is most serious for faint targets – for the same example a target near the detection limit (66 electrons total) will lose about 40% of its counts during readout. A faint background further compounds the troubles – if we had used the F502N filter instead of F555W, the last case would suffer losses of about 90% during readout.

Most CTE work has focussed on direct measurement of the photometric effects – that is, measuring target brightness as a function of distance from the readout amplifier. A new report (Biretta & Kozhurina-Platais 2005) has taken a slightly different approach and attempted to identify and quantify the charge which is left behind. They have examined apparent "tails" on hot pixels in the images, and used these to study CTE effects. Hot pixels are chosen as they are bright single pixels, and provide a simple probe of any charge which is left behind or which trails behind during readout. Figure 5 illustrates the average profile of many hot pixels from year 2002, and shows that there is a tail approximately 200 pixels long behind the hot pixel. While the surface brightness of the tail is low, it contains approximately 9% of the total charge and is equal to the expected photometric losses – i.e., this 200-pixel long tail appears to represent nearly all the photometric loss. Information about the shape and intensity of the tail should be useful in creating an image-based CTE correction.

One curious result from this study, is that there seems to be significant tails on hot pixels even near row zero on the CCD. Figure 6 illustrates the integrated intensity in the tails as a function of CCD Y position. Data are plotted for both epoch 1995 and 2001.

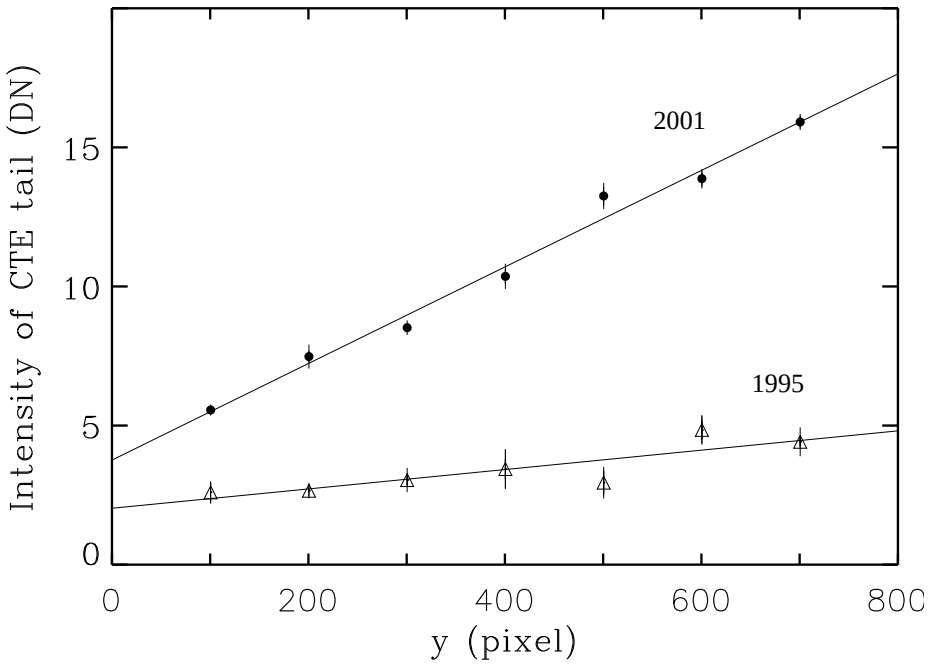

Figure 6: Measured hot pixel tail intensity (integrated) intensity as a function of CCD row at epochs 1995 and 2001. Note that the tail intensity is non-zero at $Y = 0$.

While the strength of the tails increase with Y and epoch as expected from photometric studies, the Y=0 intercept is non-zero. Apparently there is a significant amount of charge which is trapped at the position where the image is exposed, and then quickly released during the readout process. This component of photometric loss is not properly treated in the unusual CTE models which have zero loss at Y=0. Presumably there is nothing special about the Y=0 location on the CCD, and similarly charge is everywhere trapped at the position where the image is exposed. This trapping would produce an effect similar in nature and size to so-called Long vs. Short Effect (Casertano & Mutchler 1998; Whitmore & Heyer 2002) where images in short exposures are too faint relative to longer exposures of the same target.

5. Photometric Calibration

The photometric calibration of WFPC2 continues to be relatively stable. Long-term observations of the standard star GRW+70D5824 are plotted in Figure 7 for the F555W filter. There is a long-term decline in the sensitivity in all CCDs due to CTE effects, and this decrease is now approaching 4% to 5%. Aside from this long-term decrease, the photometry is stable to about 1.5% RMS.

The throughput in the UV continues to be subject to fluctuations caused by the slow build up of contaminants which are then rapidly removed during decontamination procedures. Contaminants slowly accumulate on the cold CCD windows and decrease the UV throughput. These contaminants are removed by warming the CCD assemblies to +20 °C for six hours on a regular basis. Early in the WFPC2 mission these decontaminations were carried out every 30 days, but more recent years the interval has been increased to 45 or 60 days. This increased interval is in response both to an overall reduction in contaminant levels in the camera, and also represents an effort to increase HST schedule efficiency. The UV throughput loss per day due to contaminant accumulation varies somewhat between

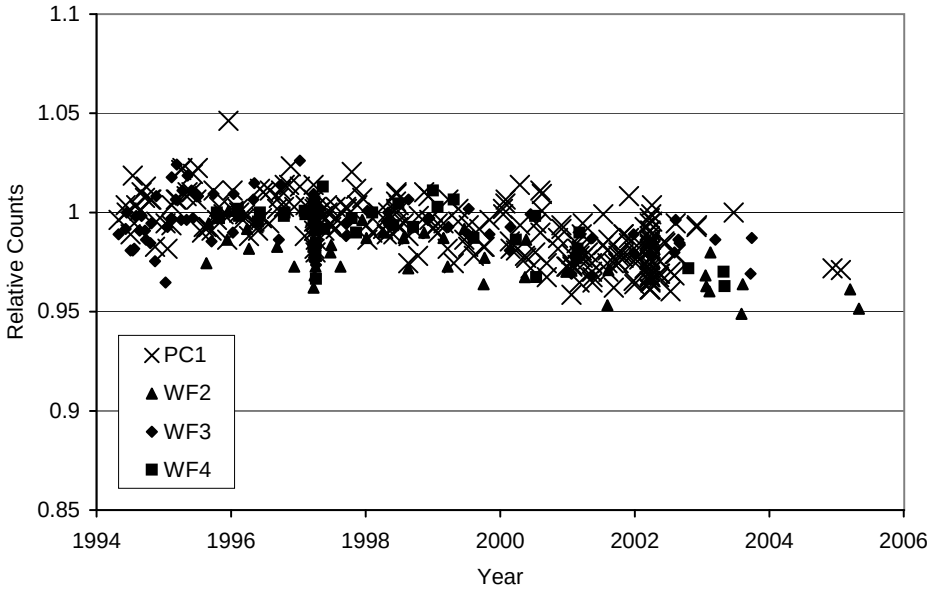

Figure 7: Relative counts observed for standard star GRW+70D5824 in filter F555W as function of time. The long-term sensitivity decrease is due to CTE effects which have not been corrected in this plot.

CCDs, but for example, in the PC1 channel early in the mission the F170W sensitivity decreased by about 0.56% per day, whereas the decrease is now about 0.32% per day.

The photometric zero points were recently reviewed by Heyer, et al., 2004. They compared zero points derived by different groups including STScI, Holtzman 1995, Dolphin 2000, Steston, and Saha. Data included both Omega Cen and NGC 2419. The zero points from different targets and different workers in the F555W and F814W filters were found to agree to about 2% RMS. But larger differences of 3% to 4% RMS were found in F439W and F336W, with the largest differences being about 6%. These differences are much larger than the Poisson errors, and reflect some unknown systematic error in the photometric calibrations.

6. Geometric Distortion

WFPC2 images suffer from mild geometric distortion which is mostly caused by refractive field flatters located immediately in front of the CCDs (these in fact act as protective windows for the CCD detectors). At intermediate wavelengths, for example in the F555W filter, the distortion reaches a total of about 6 pixels (or about 2% by pixel size) in the corners of the CCDs. As these are refractive elements, the geometric distortion will be a function of wavelength. Kozhurina-Platais, Anderson, & Koekemoer (2003) have empirically measured the wavelength dependence of the geometric distortion. They find the distortion at F300W is ∼3% greater than at F555W, and at F814W it is ∼1% less. These variations are consistent with those expected from the wavelength dependence of the refective index of MgF2. It is expected that the distortion will continue to increase into the UV. For example, the distortion at F255W is expected to be about 5% greater than that at F555W, but this has yet to be verified observationally.

7. Calibration Plans

Our current plans for Cycles 13 to 15 are to continue the routine monitors that have been established over the years (bias frames, darks, internal flats, Earth flats, standard star monitors, CTE monitors). Many of the WFPC2 filters have not had photometric calibrations in a few years, so we plan to make standard star observations in all the filters during the next year or so. We also plan to take observations of very red stars which may be useful in performing cross-calibrations between WFPC2 and ACS. No doubt many additional observations will be made to calibrate the WF4 anomaly, and in particular the low photometric results that accompany low bias levels.

8. Close-Out Calibrations

We anticipate that WFPC2 will be de-orbited near the end of 2007 during HST service mission SM4. Due to management concerns, there is also a desire to cease all WFPC2 support around that time. These two constraints together provide impetus for carefully planning our remaining WFPC2 activities, and beginning that work now. During these next two years we will embark on a number of activities to finalize and "complete" the calibration of WFPC2.

Calibration observations: obviously any and all data must be collected prior to de-orbit. An important aspect of this is to try to anticipate data which might be needed or desirable in the future. For example, data that might be useful to cross-calibrate with future instruments, or data that might be useful to resolve long-standing issues such as CTE corrections for extended sources or complex scenes.

Primary calibrations: basic calibrations should be finalized and brought up-to-date to reflect the status of WFPC2 at the end of its mission. This might potentially include new super-bias and dark reference files, new flats, and current photometric calibrations. Trends in long-term effects such as CTE should also be updated.

Documentation: final versions of the WFPC2 Instrument Handbook, Data Handbook, and various web resources must be created to capture the latest information, and ensure it will be accessible into the future.

Enhancements to WFPC2 legacy: WFPC2 will likely represent a significant fraction of the total output of HST, so it is interesting to consider projects that might enhance its value, or make the data more accessible. Some of these projects might include:

1. Image-based CTE corrections. These would properly handle complex and extended targets, and eliminate the need for future archival users to perform manual CTE photometric corrections.

2. Photometric history of WPFC2 and connection to other instruments.

3. Enhanced archival products. For example, cosmic ray rejected mosaicked images with proper photometric calibration.

There are also many areas where the calibration could be improved or made more secure. Selecting among these will obviously be a question of priorities:

1. Improved photometric zero-points, particularly in the blue and UV.

2. Better photometry for narrow-band and ramp filters. There is sometimes evidence for errors in the 5% to 10% range whereas ideally the errors should be closer to those for broad filters.

3. Narrow band filter stability: there is a potential for filter bandpasses to change over time, and this has not been addressed in any complete fashion for WFPC2.

4. Calibration of red-leaks: the blue and UV filters have significant leaks, but this has not been calibrated on-orbit in a thorough manner.

5. Sky flats for primary filters: these might provide some improvement or at least a test of the accuracy of the existing flats. Given the enormous quantity of data available, it might be possible to make reasonably noise-free sky flats.

6. Sub-pixel and focus effects on photometry: some early efforts were made, but it probably remains a significant source of uncertainty.

The HST observer community is welcomed to make suggestions for enhancements, and moreover, to actively participate in Close-Out calibration activities. Funding as well as observing time to support such work are potentially available to outside observers via the Calibration Outsource program and Calibration Proposals. We would greatly welcome outside participation in any of these projects.

References

Biretta, J., & Kozhurina-Platais, V. 2005, *Instrument Science Report* WFPC2 2005-01, available through `http://www.stsci.edu/instruments/wfpc2/wfpc2_top.html` (Baltimore: STScI)

Biretta, J., & Gonzaga, S. 2005, *Instrument Science Report* WFPC2 2005-02 (Baltimore: STScI)

Casertano, S., & Mutchler, M. 1998, *Instrument Science Report* WFPC2 1998-02 (Baltimore: STScI)

Dolphin, A. 2000, PASP, 112, 1397

Heyer, I., Richardson, M., Whitmore, B., & Lubin, L. 2004, *Instrument Science Report* WFPC2 2004-01 (Baltimore: STScI)

Heyer, I., & Biretta, J. 2005, *WFPC2 Instrument Handbook* (Baltimore: STScI)

Holtzman, J., et al. 1995, PASP, 107, 1065

Koekemoer, A. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 291

Kozhurina-Platais, V., Anderson, J., & Koekemoer, A. 2003, *Instrument Science Report* WFPC2 2003-02 (Baltimore: STScI)

Mobasher, B., et al., 2002, *HST Data Handbook* (Baltimore: STScI)

Whitmore, B. 1998, in *Proc. 1997 HST Calibration Workshop*, ed. S. Casertano et. al. (Baltimore: STScI), p. 317

Whitmore, B., Heyer, I., & Casertano, S. 1999, PASP, 111, 1559

Whitmore, B. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 281

Whitmore, B., & Heyer, I. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 359

Part 2. NICMOS

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

NICMOS Status

Roelof S. de Jong, Santiago Arribas, Elizabeth Barker, Louis E. Bergeron, Ralph C. Bohlin, Daniela Calzetti, Ilana Dashevsky, Mark Dickinson, Anton M. Koekemoer, Sangeeta Malhotra, Bahram Mobasher, Keith S. Noll, Adam G. Riess, Alfred B. Schultz, Megan L. Sosey, Thomas Wheeler, Tommy Wiklind, Chun Xu

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. We provide an overview of the most important calibration aspects of the NICMOS instrument on board of HST. We describe the performance of the instrument after the installation of the NICMOS Cooling System, and show that the behavior of the instrument has become very stable and predictable. We detail the improvements made to the NICMOS pipeline and outline plans for future developments. The derivation of the absolute photometric zero-point calibration is described in detail. Finally, we describe the discovery and quantification of a newly discovered count rate dependent non-linearity in the NICMOS cameras. This new non-linearity is distinctly different from the total count dependent non-linearity that is well known for near-infrared detectors. We show that the non-linearity has a power law behavior, with pixels with high count rates detecting slightly more flux than expected for a linear system, or vice versa, pixels with low count rate detecting slightly less than expected. The effect has a wavelength dependence with observations at the shortest wavelengths being the most affected (∼0.05-0.1 mag per dex flux change at ∼1 micron, 0.03 mag per dex at 1.6 micron).

1. Introduction

NICMOS is currently (as of 2005) the second most used science instrument on board of the Hubble Space Telescope, accounting for about 25% of its science observations. NICMOS has been operating for more than 3.5 years with the NICMOS Cooling System (NCS) that was installed in March 2002 during Servicing Mission 3B. With the NCS the instrument is operating at a very stable temperature, making it easier to calibrate than in the pre-NCS period, as many instrument characteristics show a strong temperature dependence. NICMOS has become over the run of years a more mature instrument on HST with most of its characteristics well defined and corrected in the standard data reduction pipeline. However, in recent years NICMOS has been used at the extremes of its capabilities, revealing new, unexpected instrumental effects that we are in the process of calibrating. Examples of extreme use are grism observations of the $\sim 6^{th}$ H-mag exo-solar planet host star HD209458 (PI Gilliland, ID 9642) to the Hubble Ultra Deep Field with galaxies of about 24^{th} H-mag, a dynamic range of 18 magnitudes!

The organization of the paper is as follows. In section 2 we describe the main results of the different calibration programs, in general monitoring of the instrument under normal circumstances. In section 3 we describe developments of the NICMOS calibration pipeline and plans for future improvements. Section 4 gives a detailed account of the derivation of the photometric zero-point calibration. Investigations into the recently discovered count rate dependent non-linearity are presented in section 5. Our further calibration plans are described in section 6.

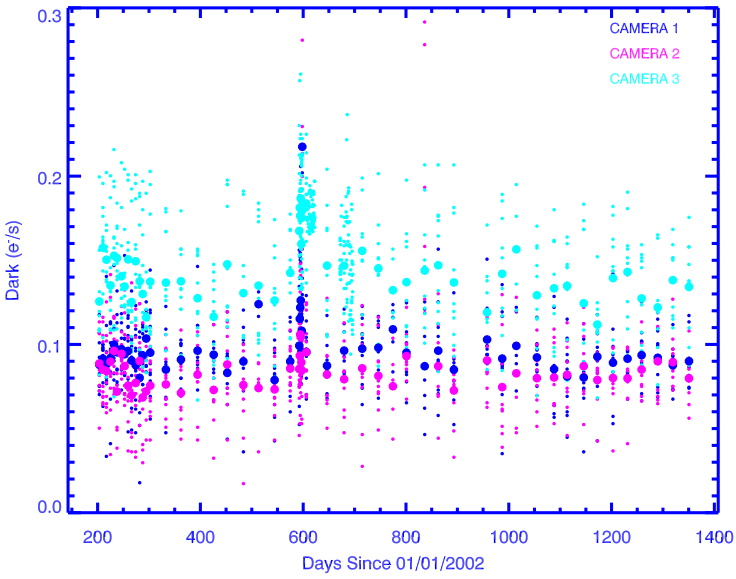

Figure 1: Dark count rates in the three different NICMOS cameras.

More detailed information on many of the NICMOS calibration efforts are described in separate contributions to these proceedings.

2. Instrument

Several programs have been executed over the lifetime of NICMOS to monitor the behavior of key calibrations. One of the most important programs is the monitoring of the instrument temperature, as many instrument characteristics depend critically on temperature. We aim to keep the temperature as close as possible to 77.1 K, adjusting the cooling rate of the NCS if necessary to account for seasonal variations and other temperature drifts.

The dark count rate is monitored at regular intervals and has been found to be stable (Fig. 1). The high count rates in this figure seen near day 600 are due to persistence of Mars observations.

The focus of the three NICMOS cameras is monitored at regular intervals using phase retrieval (Fig. 2). No focus adjustments have been necessary since the installation of NCS. NIC1 and NIC2 are permanently in focus; NIC3 remains slightly out of focus due to dewar deformations developed before NCS installation which are outside the range of corrections possible with the NICMOS focussing mechanism. No special NIC3 focus campaigns shifting the HST secondary mirror are planned, as the NIC3 camera is significantly under-sampled and little is to be gained by such a dedicated campaign.

The count rates detected in the images of the flat field monitoring program have been very stable, indicating that there has been little change in sensitivity. However, the general shape of the flat field has been changing over time since NCS installation. The deviations are now about 1-3% from minimum to maximum in NIC1 and NIC3 (NIC2 is less affected) compared to the pipeline flat fields that were created immediately after NCS installation. The effect is most severe at the shorter wavelengths. We are investigating whether these small flat field variations are the result of temperature changes using the temperature from bias method (Bergeron, in preparation). New high signal-to-noise flat field observations are being planned.

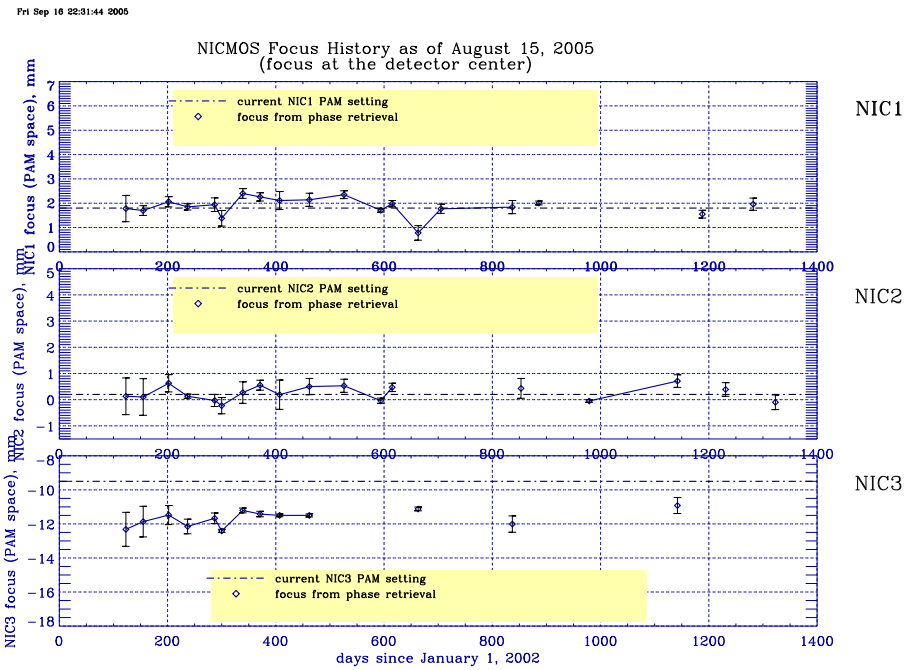

Figure 2: Focus measurements of the three NICMOS cameras. NIC1 and NIC2 are at their nominal focus setting, NIC3 remains out of reach of the PAM focussing mechanism.

A number of programs were executed to investigate NICMOS performance in 2-gyro mode (see also Sembach et al., these proceedings). The NICMOS PSF shape did not change at the resolution of the NICMOS cameras, as expected. Most critically, the coronographic rejection in the NIC2 camera did not suffer in 2-gyro mode. However, while coronographic NICMOS observations are still possible, it is no longer possible to get two roll angle coronographic observations in one orbit due to the extra overhead involved in changing guide stars under 2-gyro operations.

In June, 2005 new SPARS MULTIACCUM exposure time sequences (SAMP-SEQ) became operational, replacing the old MIF sequences. The new sequences are SPARS4, SPARS16, SPARS32 and SPARS128, complementing the already existing SPARS64 and SPARS256 sequences. These sequences have equal time steps between each readout. We recommend the use of these SPARS sequences for most observations, as they provide the most stable measurements especially in terms of amplifier glow. The alternative STEP sequences should only be used in situations where one needs to observe objects with a very large dynamic range in one observation.

Many of these calibration investigations have been described in Instrument Science Reports, which can be found at: http://www.stsci.edu/hst/nicmos/documents/isrs

3. Pipeline

Several enhancements are currently underway or have recently been implemented for NICMOS analysis routines and pipeline-related software. The MultiDrizzle software (Koekemoer et al. 2002, 2005) available within Pyraf has been extended to enable fully automated combination of calibrated NICMOS images, which can be provided either as a list of exposures or as a NICMOS association table. We have also distributed the first Pyraf release of the *SAAclean* task (Barker et al. 2005) which is based on the IDL algorithm (Bergeron and Dickinson 2003) to remove residual flux from pixels impacted by cosmic rays accumulated

during passage through the South Atlantic Anomaly. Testing is currently underway on both MultiDrizzle and SAAclean to prepare them for eventual inclusion into the automatic processing carried out by the HST archive pipeline.

Development is also currently proceeding on a task to remove the cross-talk effect known as "Mr. Staypuft", where flux from bright sources on one quadrant can be seen to propagate to pixels in corresponding locations on the other quadrants. This task will be released to the community in an upcoming Pyraf release, and may eventually also be incorporated as part of the HST pipeline. Future work will include software to determine the NICMOS detector temperature from bias and voltage measurements, as well as improved amplifier glow and pedestal correction software which will most likely make use of the improved temperature measurements.

4. Photometry zero-point

Initial absolute photometric calibrations for NICMOS were obtained during SMOV in July 1997. These preliminary results used only a few filters in each camera to establish initial corrections from the predicted ground-based vacuum measurements. These measurements obtained accuracies between 10 and 15 percent. Later observations improved the calibration to about the 5% level in all cameras and filters by May 1998. However, revised ground-based photometry from Persson et al. (1998) suggested that calibrations based on the these CDBS (Calibration Data Base System) NICMOS photometric standard stars observations might be off by as much as 5 - 14 percent in some filters, due to an 0.1 magnitude discrepancy with the ground-based photometry for the solar analog and white dwarf used for calibration.

With the installation of the NCS the NICMOS detectors are operating at dramatically higher temperature with an associated strong change in sensitivity. Furthermore, the limited accuracy of previous analyses, new reduction routines, improvements to the pipeline calibrations and calibration images, and a better understanding of the instrumental anomalies encouraged a complete re-analysis of all available NICMOS photometric calibration data. The following paragraphs will discuss calculation of new aperture corrections for each camera/filter combination, improvements of the spectrophotometric standard star spectra used for calibration and an assessment of the overall photometric stability during the full lifetime of the instrument.

All standard star observations were reduced with the latest version of the calibration pipeline and reference files. In post-pipeline processing, quadrant dependent bias was removed with *pedsky* and residual readout shading was removed by 1D fitting of the sky.

4.1. Aperture Corrections

Previous calibrations used a single aperture correction regardless of the filter element to correct the fixed aperture standard star photometry to an infinite aperture. While this can be used consistently for point source observations, the wavelength dependence of the PSF makes this incorrect for calibrating extended sources and makes comparisons with ground-based observations difficult. New, wavelength dependent, aperture corrections for each of the cameras have been calculated using the TinyTim PSF modeling software. Figure 3 shows the new aperture corrections compared to the previous fixed corrections. The new aperture corrections, along with the aperture sizes and sky annuli used, are tabulated on the NICMOS photometry web site.

4.2. Absolute Spectrophotometric Standards

The photometric calibration keywords are derived from the comparison of the measured NICMOS count rates for the standard star observations to the spectral flux density of the standard star averaged over the NICMOS bandpass. There are no ground-based spectrophotometric observations of standard stars with complete coverage over the NICMOS

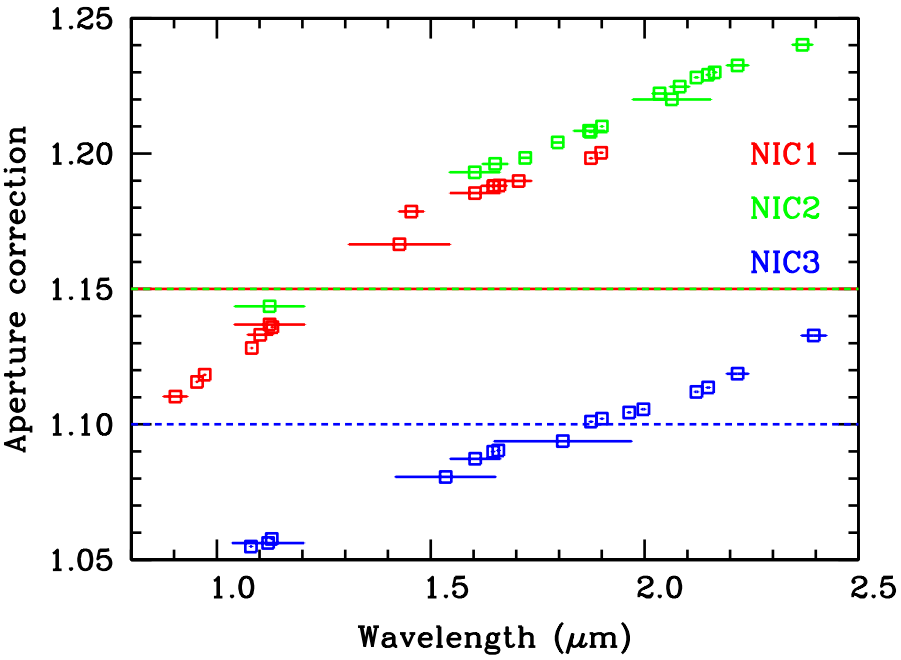

Figure 3: Comparison of the old aperture corrections (dashed lines) with the new aperture corrections (squares, the filter widths are indicated by the horizontal lines).

wavelength range, and therefore we must use accurately calibrated "surrogate spectra" instead in our comparison. This was the motivation for using solar analog and white dwarf standards for the NICMOS photometric calibration. The absolute spectral energy distribution of the Sun is well known (see Colina, Bohlin & Castelli 1996 and references therein), and thus can be scaled reliably to represent the spectrophotometry for solar analog standard stars like P330E. DA white dwarfs like G191B2B have relatively simple stellar atmospheres, and considerable effort has gone into accurately modeling these and comparing them to UV-through-optical spectrophotometry (Bohlin, Colina & Finley 1995; Bohlin 1996; Bohlin 2000).

As described by Colina & Bohlin (1997), the infrared spectrum of P330E is represented by the solar spectrum from Colina, Bohlin & Castelli (1996). The white dwarf G191B2B is represented by an LTE model calculated by D. Finley (described in Bohlin 2000). These spectrophotometric models are then normalized using ground-based photometry of the NICMOS standard stars. Persson et al. (1998) have obtained ground-based JHK photometry for a large set of faint infrared standard stars, including the HST/NICMOS solar analog standards. Persson (private communication) also observed G191B2B as part of the same program.

In order to normalize the standard star spectral models, we must convert Persson's JHK magnitudes to absolute flux density units. Campins, Rieke & Lebofsky (1985) provide an absolute infrared flux calibration scale using a solar analog method. However, the effective wavelengths and bandwidths of their JHK filters (which we will refer to as the Arizona system, where Vega is defined to have $J{=}H{=}K{=}0.02$) differ somewhat from those used by Persson et al. (calibrated to the CIT system, where Vega is defined to have $J{=}H{=}K{=}0.0$). In order to shift the Campins et al. absolute calibration to the Persson et al. bandpasses, we have used an ATLAS 9 atmosphere model for Vega. This model is not used for any absolute calibration, but simply to compute flux density ratios for Vega between the Arizona and Persson et al. bandpasses. These are then used to convert the Campins et al. Vega flux densities to the Persson et al. bandpasses, and hence to provide the absolute flux density

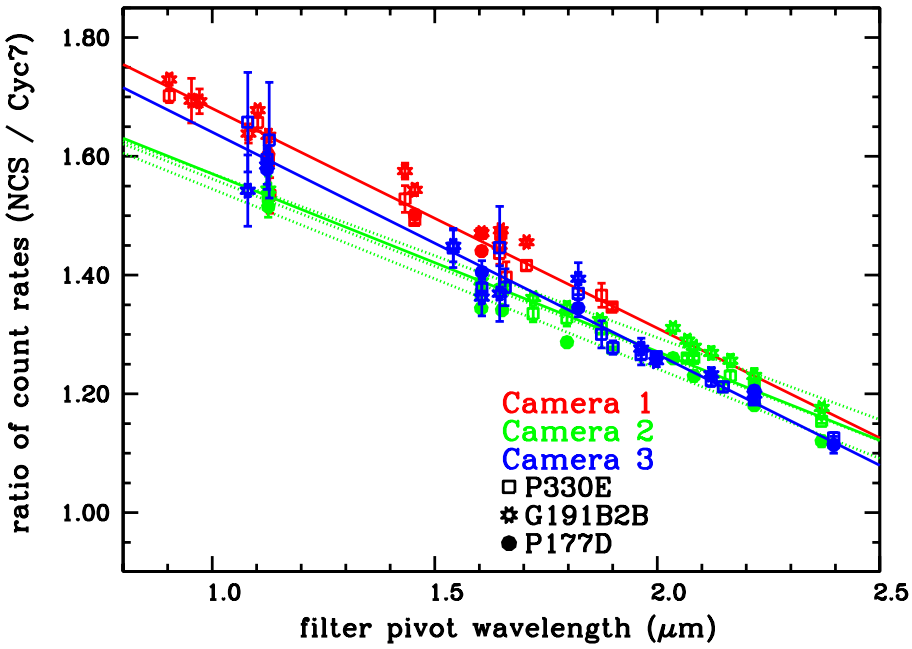

Figure 4: The relative change in count rate of standard stars from Cycle 7 to post-NCS observations. The lines are linear fits to the data of the three cameras.

calibration for the Persson et al. measurements. In this way, m=0 is calibrated to be 1626, 1056, and 658 Jy for the Persson et al. *JHK* bandpasses, respectively.

The P330E and G191B2B spectrophotometric models are then synthetically integrated through the Persson et al. *JHK* passbands, and the bandpass-averaged flux densities are converted to magnitudes for comparison to the ground-based photometry. This comparison indicates that the Colina & Bohlin (1997) P330E model requires an average flux renormalization of +7% to match the *JHK* photometry at the 0.01 mag level. For G191B2B, the synthetic and ground-based photometry agree precisely (0.002 mag) at J and within 0.02 mags at K, but differ by 0.053 mags at H. We adopt the Bohlin (2000) G191B2B model without change, but note this possible discrepancy near 1.6 microns.

4.3. Zero-point Calibration

For each standard star we determined the average aperture count rate for all NIC1, NIC2, and NIC3 filters for each dither position and the many repeat observations. Observations near bad pixels or other outliers were removed from the average. The observed, aperture corrected, average total count rates for each star in each filter were compared with the predicted total count rates obtained using SYNPHOT synthetic photometry package[1] on the spectra of G191B2B and P330E. We used the *calcphot* program to calculate the effective stimulus of a source with a flat spectrum in f_ν. We used the ground-based filter shapes and detector quantum efficiency QE curve for the Cycle 7 data. The QE curve was modified for the post-NCS change in sensitivity by multiplying the QE with a linear correction that was determined from comparing pre-NCS to post-NCS count rates of standard stars (Fig. 4). Comparing the thus calculated effective stimulus to the observed average count rates results in the PHOTFNU keywords. Similar calculations provide the PHOTFLAM, PHOTPLAM and PHOTBW values found in NICMOS image headers, as described in the

[1]available at http://www.stsci.edu/resources/software_hardware/stsdas/synphot

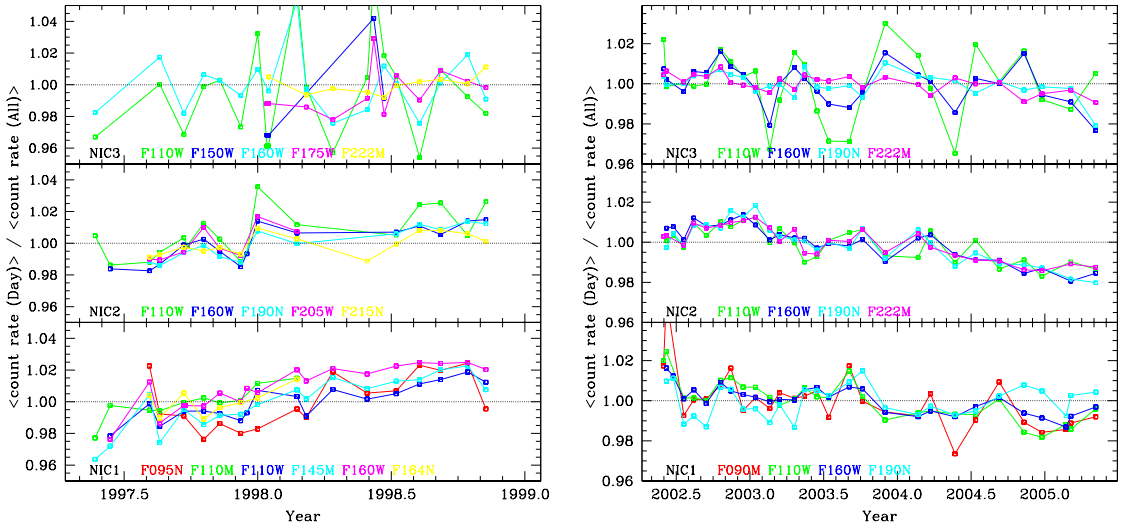

Figure 5: Evolution of the P330E standard star count rates: (left) the Cycle 7 data; (right) the post-NCS data.

SYNPHOT manual and in Sirianni et al. (2005). The new calibration values have been automatically provided in the image headers retrieved from the HST archive since June 2004. The new filter throughput files were made available in the CDBS in December 2005. The latest calibration values can also be found at the NICMOS photometry web pages: http://www.stsci.edu/hst/nicmos/performance/photometry

4.4. Photometric Stability

Photometric stability of NICMOS was investigated using all observations of Solar analog star P330E used in the photometric monitoring program (Fig. 5). The Cycle 7 data show the clear increase in sensitivity due to the increase in temperature while the solid nitrogen evaporated. After temperature correction there is a slight decrease in sensitivity. The post-NCS data show clear decrease in sensitivity in NIC2, a downward trend in NIC1, but the NIC3 data is too noisy due to intra-pixel sensitivity variations to tell whether there is any downward trend. The cause of this decrease in standard star sensitivity is not clear yet, as it is not matched in the flat field lamp monitoring data. However, similar downward trends are seen in the few repeat observations of G191B2B.

5. Non-linearity

In a recent analysis of NICMOS, STIS and ACS spectral data Bohlin et al. (2005) found that NICMOS shows a systematic count rate dependent non-linearity, primarily at the shorter wavelengths (Fig. 6). The same spectra show a similar non-linearity when compared to ACS photometry. This count rate dependent non-linearity is distinctly different from the normal non-linearity of near-infrared detectors that depends on the total counts, not on the count rate. The total count non-linearity is well understood and corrected in the NICMOS pipeline.

The non-linearity is such that at high count rates there are more counts than expected, and at low count rates less than expected, compared to intermediate count rates. The non-linearity shows no sign of turnover over the full 4 orders of magnitude measured (Fig. 6), hence we cannot say with confidence that we are missing photon detections at the faint end

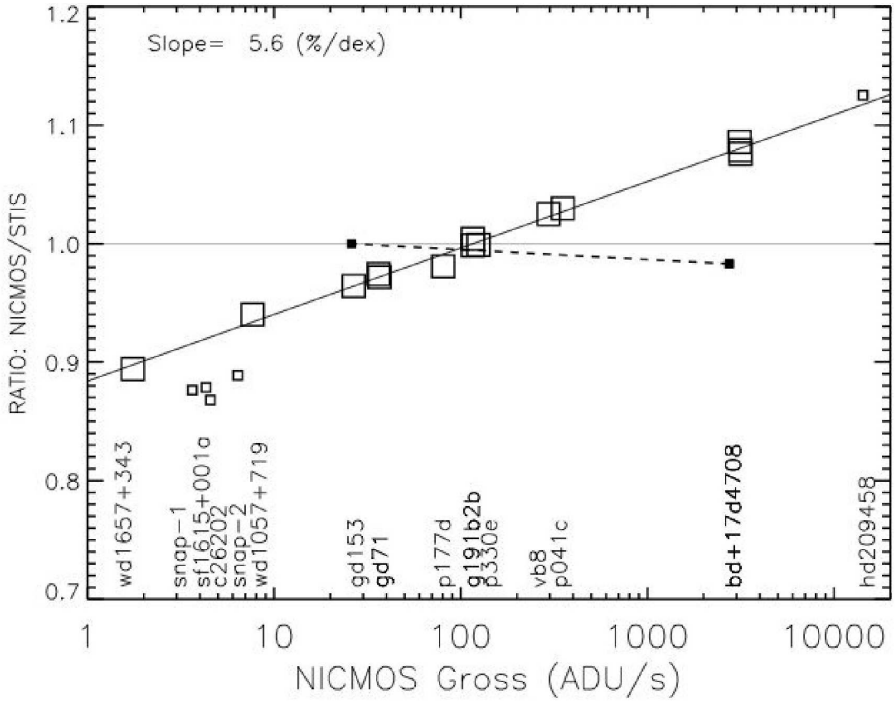

Figure 6: Ratio of NICMOS grism fluxes to the STIS measurements averaged over the 0.82–0.97μm overlap range, normalized to the P330E ratio (big open squares) and least square fit (solid line). There is a non-linearity of about 5.6% per dex or a total of 23% over the 4 dex dynamic range in observed response. Small open squares—Observations not used for assessing the linearity. Of the six faintest stars, only the two WD stars with good models are used for the fit, because the STIS fluxes seem to be systematically overcorrected for CTE by a few percent. The bright star HD209458 was observed with NICMOS in a defocused mode. Small filled squares connected by a dashed line—ACS grism fluxes compared to STIS in the same 0.82–0.97μm band. Over a dynamic range of $\sim$100$\times$ between GD153 and BD+174708, the CCD detectors on STIS and the HRC in ACS measure the same relative flux to within 2%.

or are getting extra detections at the bright end. The non-linearity is well modeled by a power law.

The count rates of the NICMOS spectra show in general good agreement with the NICMOS photometry count rates of the same objects, indicating that the NICMOS system is internally consistent and that it is not the spectral data reduction that is at fault. A few more indications have been found that NICMOS suffers from a non-linearity dependent on the incoming flux: 1) narrowband filters at the shorter wavelengths required larger in-flight corrections from their ground-based determined throughputs than the broadband filters, 2) high redshift supernova fluxes are slightly fainter in F110W than expected based on their ACS fluxes and well tested SN models (Adam Riess, private communication), and 3) galaxies in the HUDF are slightly fainter than expected based on ACS and ground-based J&K magnitudes combined with SED modeling (Mobasher & Riess 2005; Coe et al., these proceedings).

However, all these lines of evidence rely on modeling of filter throughputs and/or spectral modeling of sources. Here we describe a test that depends on the change in incoming flux on the detector alone. NICMOS is a shutterless instrument and observes the sky while obtaining calibration flat fields using its internal lamps. The same object can be observed with an artificially increased flux and the count rate for objects can be compared with and

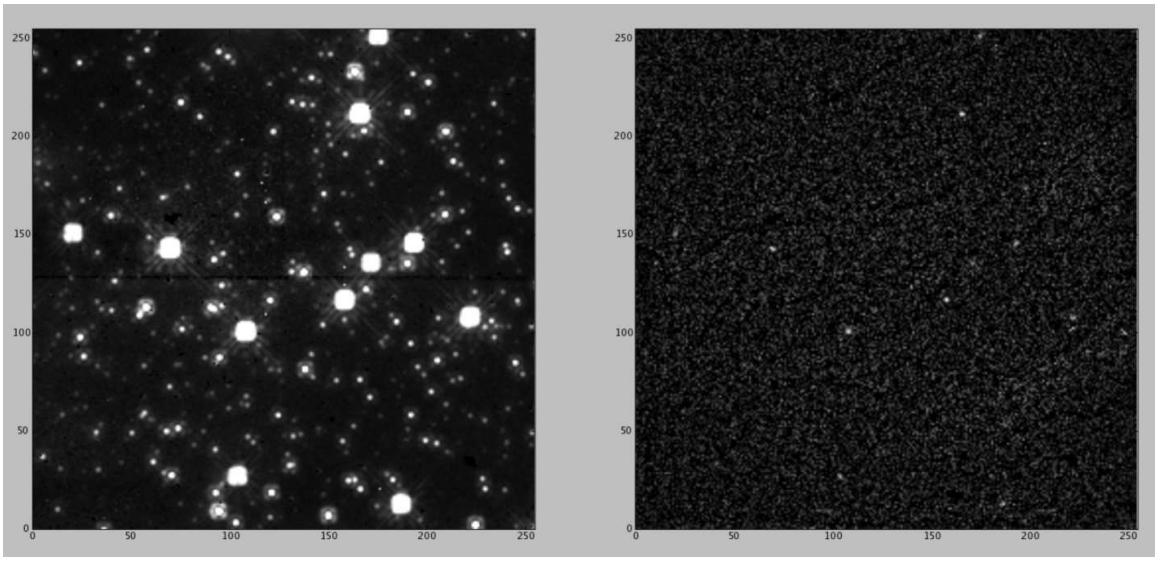

Figure 7: (Left) NIC1 F110W lamp-off image of NGC 1850. (Right) NIC1 F110W lamp-on minus lamp-off image. Bright stars are clearly not well subtracted and leave residual flux as expected for count rate dependent non-linearity.

without extra lamplight. For a fully linear system, adding a background flux should not enhance the flux in the object, but any flux-dependent non-linearity is revealed immediately when subtracting lamp-off images from lamp-on images.

Star cluster NGC 1850 was observed in the Cycle 14 calibration program in a sequence of lamp-off, lamp-on, and lamp-off using the same telescope pointing and without changing the exposure sequence. Similar observations taken in Cycle 7 for a different purpose were also analyzed. The data were analyzed under the assumption that a power law can model the non-linearity:

$$cr(x,y) \propto (f_{tot}(x,y))^{\alpha},$$

with $cr(x,y)$ the measured count rate in ADU/s and $f_{tot}(x,y)$ the total flux falling on a detector pixel at (x,y). For a non-linearity of $\sim$5% per dex this corresponds to $\alpha\sim$1.02. In magnitudes we have an offset of $\Delta m = 2.5(\alpha - 1)$ per dex change in incident flux. When we subtract the lamp-off from the lamp-on observation we expect to see positive residuals at positions where there are objects if α>1:

$$cr_{on} - cr_{off} \propto (f_{obj} + f_{sky} + f_{lamp})^{\alpha} - (f_{obj} + f_{sky})^{\alpha} \sim (f_{obj} + f_{lamp})^{\alpha} - (f_{obj})^{\alpha},$$

where it is assumed that the sky flux is small compared to the other fluxes. Such image residuals are shown in Figure 7. The absolute boost in measured count rate is largest for bright objects, but the relative increase in measured count rate is larger for lower object fluxes, because the relative increase in flux by switching on the lamp is much larger. However, at low count rates the noise dramatically increases and we have to average many points to see the effect. This is shown in Figure 8, where we plot both the absolute and relative count rate increase due to the lamp background.

The fitted non-linearity functions are overplotted in Figure 8 and the measured α values are tabulated in Table 1. A number of points can immediately be taken from the Table. NICMOS has a significant count rate dependent non-linearity, becoming more severe at shorter wavelengths. This is a different non-linearity from the well-known total count dependent non-linearity. The non-linearity in NIC1 and NIC2 amounts to 0.06-0.10 mag offset per dex change in incident flux for the shortest wavelength (F090M and F110W), about

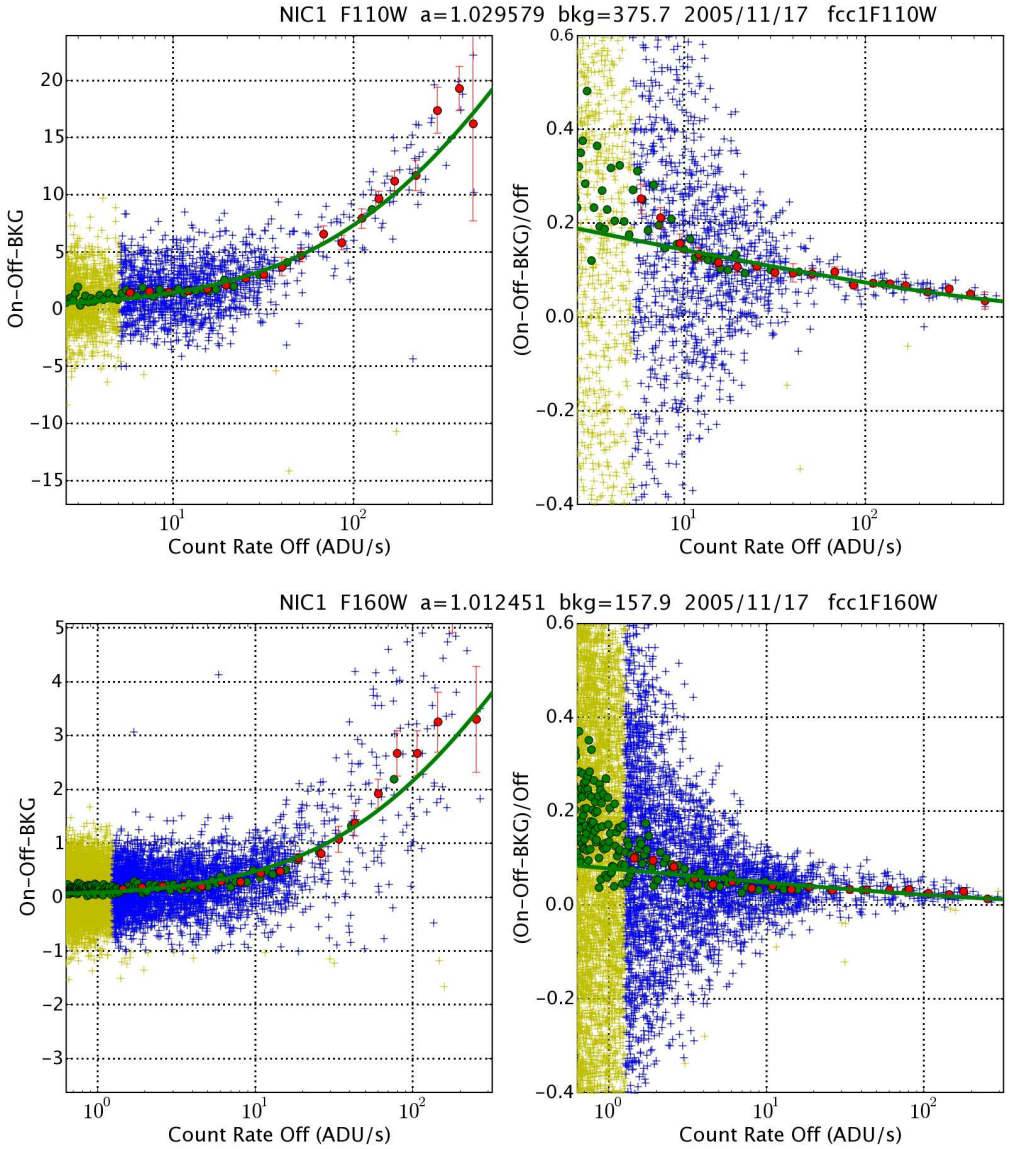

Figure 8: The absolute and relative difference in lamp-on minus lamp-off count rates as function of the lamp-off count rates on a pixel-by-pixel basis. The yellow + symbols are for all data, the blue + symbols for the data used in the fit (bad and low S/N points filtered). The green circles are the averages in bins of 50 pixels in ascending lamp-off count rate for all pixels. The red circles are binned averages in 30 equal logarithmic steps in lamp-off count rate for the selected pixels. The green lines are the fitted non-linearity functions, with the fitted α parameter labeled at the top. Left is the absolute difference ($cr_{on} - cr_{off} - cr_{lamp}$), right is the relative difference ($cr_{on} - cr_{off} - cr_{lamp}$)/ cr_{off}. While the brightest points have the largest flux change in absolute sense and are easy to measure above the noise (and are not due to background subtraction errors), the fainter points change relatively the most and have larger calibration errors relative to bright standard stars. Top) NIC1 camera, F110W filter. Bottom) NIC1, F160W filter.

Date	Camera	Filter	α	Δm/dex
		Cycle 7		
1998/02/18	2	F110W	1.022 ± 0.001	0.055 ± 0.003
1998/04/17	2	F110W	1.025 ± 0.004	0.063 ± 0.010
1998/06/04	2	F110W	1.023 ± 0.001	0.059 ± 0.002
1998/08/06	2	F110W	1.024 ± 0.001	0.061 ± 0.002
1998/09/24	2	F110W	1.022 ± 0.001	0.054 ± 0.002
		Cycle 14		
2005/11/17	1	F090M	1.040 ± 0.003	0.101 ± 0.008
2005/11/17	1	F110W	1.030 ± 0.003	0.074 ± 0.009
2005/11/17	1	F160W	1.012 ± 0.002	0.031 ± 0.006
2005/11/17	2	F110W	1.025 ± 0.002	0.063 ± 0.006
2005/11/17	2	F160W	1.012 ± 0.006	0.029 ± 0.015
2005/11/17	2	F187W	1.005 ± 0.004	0.013 ± 0.009

Table 1: Measured α and Δm values in Cycles 7 and 14.

0.03 mag/dex at F160W and less than that at longer wavelengths. These corrections are larger than predicted from the Bohlin et al. (2005) NIC3 grism results, which may point to intrinsic detector differences or might be the result of a different analysis method. The non-linearity seems to have changed very little from Cycle 7 to Cycle 14 (in F110W NIC2), and hence is unlikely to depend on detector temperature. The fact that there is a wavelength dependence to the effect in the lamp off/on/off test and that this trend quantitatively agrees with the grism observations strongly argues against this being the result of a data reduction error and that the cause is intrinsic to the measurement.

To what extent NICMOS photometry is affected by the non-linearity depends on the wavelength of the observations (i.e. the α parameter), whether the object is a point source or extended, and on the count rate of the sky background (as the count rate will never go below the sky level and hence the non-linearity will level off, even if the sources have lower count rates). Given that the NICMOS standard stars are of about the 12th magnitude, the maximum expected offset for the Hubble UDF for example is about 0.15-0.2 mag at 22 F110W AB-mag, where the objects are comparable to or below the sky count level (see contribution by Mobasher, Thompson and Coe in these proceedings for further analysis of the HUDF).

The NICMOS team has as of yet not found a physical explanation for the count rate non-linearity. However, with these measurements we can start to calculate corrections to point source photometry. Software is being developed for more complicated cases with extended sources or objects close to the sky background level. Eventually corrections will be incorporated in the NICMOS calibration pipeline. Non-linearity lamp on/off calibration observations for all regularly used filter/camera combinations are currently being planned.

6. Calibration plans

The NICMOS team is currently implementing a calibration plan that addresses several of the anomalies observed above. New, high signal-to-noise flat field observations will be obtained for all filters regularly used. Dither observations with NIC1 have been obtained of star cluster NGC 1850 to investigate the low frequency flat field variations seen in the photometric monitoring data. Lamp on/off observations of star fields will be obtained in all regularly used camera/filter combinations to investigate the count rate dependent non-linearity. Details of the NICMOS calibration plan for Cycle 14 can be found in the contribution by Arribas to these proceedings.

References

Barker, E. A., Koekemoer, A. M. and Laidler, V. 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 175

Bergeron, L. E. and Dickinson, M. 2003, *Instrument Science Report* NICMOS 2003-010 (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Bohlin, R. C. 1996, AJ, 111, 1743

Bohlin, R. C. 2000, AJ, 120, 437

Bohlin, R. C., Colina, L., & Finley, D. S. 1995, AJ, 110, 1316

Bohlin, R. C., Lindler, D., Riess, A. 2005, *Instrument Science Report* NICMOS 2005-002 (Baltimore: STScI)

Campins, H., Rieke, G. H., Lebofsky, M. J. 1985, AJ 90, 896

Colina, L., & Bohlin, R. 1997, AJ, 113, 1138

Colina, L., Bohlin, R. C., & Castelli, F. 1996, AJ, 112, 307

Koekemoer, A. M., Fruchter, A. S., Hook., R. N., & Hack, R. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI)

Koekemoer, A. M., Fruchter, A. S., Hook., R. N., Hack, R., & Hanley, C. 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 423

Mobasher, M., & Riess, A. 2005, *Instrument Science Report* NICMOS 2005-003 (Baltimore: STScI)

Persson, S. E., Murphy, D. C., Krzeminski, W., Roth, M., & Rieke, M. J. 1998, AJ, 116, 2475

Sembach, K. et al. 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 375

Sirianni, M., et al. 2005, PASP, 117, 1049

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

NICMOS Calibration Plans for Cycles 13 and 14

Santiago Arribas, L. Eddie Bergeron, Roelof de Jong, Anton Koekemoer, Sangeeta Malhotra, Bahram Mobasher, Keith Noll, Al Schultz, Tommy Wiklind, Chun Xu

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract.

This paper summarizes the NICMOS Calibration Plans for Cycles 13 and 14. These plans complement the SMOV3b, the Cycle 10 (interim), and the Cycles 11 and 12 (regular) calibration programs executed after the installation of the NICMOS Cooling System (NCS). The special calibrations on Cycle 13 were focussed on a follow up of the spectroscopic recalibration initiated in Cycle 12. This program led to the discovery of a possible count rate non-linearity, which has triggered a special program for Cycle 13 and a number of subsequent tests and calibrations during Cycle 14. At the time of writing this is a very active area of research. We also briefly comment on other calibrations defined to address other specific issues like: the autoreset test, the SPARS sequence tests, and the low-frequency flat residuals for NIC1. The calibration programs for the 2-Gyro campaigns are not included here, since they have been described elsewhere (see Sembach et al. 2006, in these proceedings). Further details and updates on specific programs can be found via the NICMOS web site.

1. Introduction

Since the installation of the NICMOS Cooling System (NCS) during the Servicing Mission 3b, NICMOS has shown a very stable behavior as a consequence of its well controlled operating temperature (Wiklind & Wheeler, 2006). This stable behavior has had an important impact on the calibration plans, which have evolved into three main directions. First, the frequency of some of the monitoring programs have been gradually reduced. Second, due to the instrument's improved stability, the systematic errors in the calibration have become relatively more important. Consequently a reduction in the systematic errors in the zero points and grism sensitivity curves has been pursued. Third, we have also studied the behavior of the instrument at extreme count rate regimes (e.g. UDF, very high S/N spectrophotometry). Apart from these generic lines, other calibrations to address specific problems (e.g. autoreset, new multi-accum sequences, etc.) have been also implemented. In the following paragraphs we will briefly comment on each of these categories, pointing out the most relevant aspects of the Cycles 13 and 14 Calibration Plans.

1.1. Monitor programs

As a consequence of the stable behavior shown by the instrument, the frequency of the different monitoring programs, which track the stability of key properties of the instrument (i.e. darks, flats, focus, photometry), have been gradually reduced between Cycles 10 and 13. Although the possibility was discussed of reducing the frequency of some programs even further by combining their information (e.g. photometric stability could, in principle, be checked using the photometric monitoring data), it was considered important to maintain homogeneity in the programs across cycles. During Cycle 13 the major reductions affected the flat and focus programs, whose frequency were cut by factors of 2 and 1.5, respectively. During Cycle 14 we maintained the same frequency as in Cycle 13 for all monitor programs.

However, the change to 2 Gyro operating mode for HST has reduced the visibility period of some of the objects traditionally used for some of these programs, like the photometric and focus monitor. Consequently we have added new objects to these programs (G191B2B as a new photometric standard and 47Tuc as a new focussing cluster).

1.2. Photometric and Spectrophotometric recalibrations

The more stable behavior of the detectors has led to the systematic errors playing a relatively more important role in limiting the actual accuracy of standard observations. Consequently we launched in Cycle 12 two programs (photometric and spectrophotometric re-calibrations) aimed at improving the systematic errors associated with the photometric zero-points and sensitivity curves with respect to those achieved during Cycle 7 and 7N. The photometric re-calibration expanded the number and type of stars used for obtaining the zero points for all filters. Similarly, the spectroscopic recalibration allowed a reanalysis of the sensitivity curves for the three grisms using a larger number and variety of stars. Several of the selected objects were common for the two proposals. Therefore, these data allowed a spectrophotometric recalibration by cross-checking results obtained independently from these modes (i.e. imaging/filters and spectroscopy/grism). This program (which included new faint calibrating stars) has led to the discovery of a possible count rate non-linearity. This effect, the so called 'Bohlin effect', is described in detail in Bohlin, Lindler, & Riess (2005). As a response to this effect, a special calibration program was implemented in Cycle 13, and several tests and calibrations have followed during Cycle 14. This effect is dependent on the count rate and thus affects observations in a source-dependent way, as we briefly describe below. The new data obtained with the photometric re-calibration program (as well as the larger set of data collected with the photometric monitoring programs) also suggested the possibility of some systematic residual associated with the position of the object in the field of view. During Cycles 7 and 7N the possible variation of the response (after flat-fielding) as a function of the position in the field of view was estimated to be small compared with the actual uncertainties in the photometry. However, with the current higher precision data this is unclear, and a study to analyze the low frequency flat residuals has been implemented in the Cycle 14 Calibration Plan. This will be done initially for camera 1 since the intra-pixel sensitivity effect on the photometry is smaller. Depending on the results of this program a similar one could be also implemented for cameras 2 and 3 (which require extensive dithering to average out this effect; Xu & Mobasher, 2003).

1.3. Calibrating the extremes: From the very high S/N regime to the UDF

During Cycles 12-14 the behavior of NICMOS at extreme count rate regimes has been investigated. On one hand a Cycle 12 calibration proposal lead by R. Gilliland was aimed at investigating the relative spectrophotometric accuracy of the instrument in the very high S/N regime. It was found that NICMOS can reach a relative precision of around 10^{-4} or even higher in time series observations (Gilliland and Arribas 2003). Further details can be found in Gilliland (2006). In the other extreme, the Hubble Ultra Deep Field observations imposed new calibration challenges. Details on these particular calibrations are discussed in Mobasher & Riess (2005) and Thompson (2006). Because objects in the UDF are observed at extremely low count rates (i.e. very different to those typical of standard stars observations), their calibration may be affected by the 'Bohlin effect' mentioned above. In order to understand better the origin of this reported non-linearity, a program which includes several tests has been implemented in the Cycle 14 Calibration Plan. These tests are: i) the measurement of the wavelength dependence of the persistence decay after an exposure of a bright star in a series of multi-accum dark frames, ii) the measure of the non-linearity dependence on the count rate by observing a field of stars in a sequence of lamp off/on, and so artificially increasing the background level (this will be done in imaging and grism modes), and iii) the photometric measurements of the faint standard stars

SNAP-2 and WD1657+343 (on which the NICMOS non-linearity was originally discovered using grism observations) will be repeated to obtain higher S/N. All these observations should clarify the origin of the non-linearity effect, and its possible relation with the UDF calibrations (de Jong 2006).

1.4. Specific calibrations

Other programs implemented in the calibration plans addressed more specific issues like:

i) *Delta-T program*: The goal of this program was to study the effects of the detector temperature on the darks. (This was a Cycle 12 calibration program implemented too late to be included in the standard calibration plan for this cycle).

ii) *Autoreset test*: This Cycle 13 calibration program is aimed at testing the short term temperature ripples in the NICMOS dewar.

iii) *New SPARS sequences*: This Cycle 14 program had the goal of testing the darks of the newly implemented SPARS multiaccum sequences.

Further details on the individual proposals may be found in the table below (see also Arribas et al. 2005), and via the standard HST and NICMOS web sites.

Acknowledgments. Thanks are due to Ralph Bohlin, Ron Gilliland, Dean Hines, Marcia Rieke, Adam Riess, Glenn Schneider and Rodger Thompson for their contributions to the recalibration of NICMOS after the installation of NCS.

References

Arribas, S. et al.. 2005, *Instrument Science Report* NICMOS 2005-005 (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Bohlin, R., Lindler, R., & Riess, A. 2005, *Instrument Science Report* NICMOS 2005-002 (Baltimore: STScI)

Gilliland, R., & Arribas, S. 2003, *Instrument Science Report* NICMOS 2005-001 (Baltimore: STScI)

Gilliland, R. 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 162

de Jong, R. 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 121

Mobasher, B., & Riess, A. 2005, *Instrument Science Report* NICMOS 2005-003 (Baltimore: STScI)

Sembach, K. et al. 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 375

Thompson, R. 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 137

Wiklind, T, & Wheeler, T. 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 169

Xu, C., & Mobasher, B. 2003, *Instrument Science Report* NICMOS 2003-009 (Baltimore: STScI)

Table 1: Cycle 13 and Cycle 14 calibration programs. Details on individual programs can be obtained via the HST-STScI web site at http://www.stsci.edu/hst

Activity title	ID (Cycle/Program)	Comments
Multiaccum Darks	10380 (C13), 10723(C14)	Monitor programs (monthly). Linear component of the dark current, shading, ampified glow.
Focus Stability	10382 (C13), 10724 (C14)	Monitor programs. NIC1 and NIC2 every 45 days. A new field (47Tuc) added in Cycle 14 to account for 2Gyro operations.
Photometric Stability	10381 (C13), 10725 (C14)	Monitor programs (montly visits). Same as in previous cycles (i.e. observations of P3003E with selected broad filters). A new object added in Cycle 14 to account for 2Gyro operations.
Flat Fields Stability	10379 (C13), 10728 (C14)	Monitor programs. Pointed lamp observations using selected filters.
Grism Recalibration: follow up	10383 (C13)	Special calibration program. Observations of the standard star BD+17 4708. Led to the discovery of a possible count rate non-linearity.
Extreme count rates linearity test	10454 (C13)	Special calibration program. First response to the discovery of a possible count rate non-linearity.
Auto-reset test	10465 (C13)	Special calibration program to test the short term temperature ripples in the NICMOS dewar.
Non-linearity tests	10726 (C14)	Special calibration program. Includes several tests to analyze the count rate dependent non-linearity seen in NICMOS spectrophotometry.
Low-frequency flat residuals for NIC1	10727 (C14)	Special calibration program to better characterize the flat-field errors seen for camera 1 in the photometric monitoring data.
Test for newly implemented SPARS sequences	10721 (C14)	Special calibration program to measure the dark current of the new multiaccum sequences for the three NICMOS detectors.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

NICMOS Calibration Challenges in the Ultra Deep Field

Rodger I. Thompson

Steward Observatory, University of Arizona, Tucson, AZ 85721

Abstract. The reduction of NICMOS observations in the Hubble Ultra Deep Field required unique reduction and calibration techniques, not required for images containing brighter sources. This paper describes some of these techniques which may be useful in the reduction of other fields containing extremely faint sources. These techniques apply to NICMOS camera 3 which was the only NICMOS camera utilized in the observations. The absolute accuracy of the NICMOS camera 3 photometry was also investigated using the observations of the solar analog star P330-E. As a result the sensitivity of the camera in the F110W and F160W was found to be less than the sensitivities used in the construction of the Version 2.0 NICMOS Treasury catalog for the observations in the Hubble Ultra Deep Field. The catalog fluxes are too low by between 8 and 9%.

1. Introduction

In conjunction with the Director's Discretionary Time Advanced Camera for Surveys (ACS) observations of the Hubble Ultra Deep Field (HUDF) a Treasury Program of NICMOS images was conducted using the F110W and F160W filters in Camera 3. A detailed description of the data reduction and photometry is given in Thompson et al. (2005). This paper will elaborate on some procedures and contains additional calibrations obtained after the publication of Thompson et al. (2005). It is expected that the current treasury products (Version 2.0) will be amended to reflect the new calibrations. The treasury high level science products and processed images are available from the Multimission Archive at STScI (MAST). Due to the smaller size of the NICMOS camera 3 field of view only a subset of the ACS HUDF image was covered. We will refer to that field as the NICMOS Hubble Ultra Deep Field (NHUDF).

2. Observations

The NHUDF covers 5.4 sq. arc minutes roughly in the center of the HUDF. Equal observing time was allocated to the F110W and F160W filters. Each orbit produced one 1344 second integration in each of the filters. The F110W observation always preceded the F160W integration. The orientation was chosen to be as close to the orientation of the ACS images as possible. The observing period was divided into two epochs with the NICMOS orientations clocked by 90 degrees between the two epochs. The STScI scheduler was able to make all orbits free of SAA passages which greatly simplified the data reduction. The images were taken in a 3x3 grid and dithered to reduce the effects of intrapixel sensitivity variations. The final F110W and F160W images supplied to MAST are 144 arc seconds on each side. See Thompson et al. (2005) for further details.

3. Photometric Challenges of the NHUDF

Although the NICMOS images did not reach the depth of the ACS images due to the reduced observing time per unit area, they did reach a level that required special reduction techniques. To characterize the sensitivity the 1σ noise flux is 4.6×10^{-9} Jy per pixel which is equivalent to 1 detected photon every 450 seconds which is an average of 3 detected photons per integration. One analog to digital converter unit (ADU) is 6.5 electrons therefore the 1σ level is 0.5 ADUs per integration. These signals are superimposed on a sky signal due to zodiacal light of 1 detected photon per second.

4. Dark Images

The NICMOS dark current images have very significant structure, often referred to as shading. This structure has an amplitude that is much larger than the signals due to faint galaxies, therefore, proper subtraction of accurate dark images is essential for proper reduction of the NHUDF images. The dark images are also temperature dependent so that even small changes in the detector temperature can alter the dark structure. To ensure proper dark images, a dark image was taken in each NHUDF orbit during the occultation period of the orbit. The dark images were then medianed to produce an accurate dark image that was contemporaneous with the science images for each read in the SPARS64 multiaccum readout sequence.

There was some concern about the darks since the temperature set point of the NICMOS Cooling System (NCS) was changed during the course of the observations to adjust for seasonal changes of the heat load on HST. Examination of the darks previous to the set point change and after the change did not reveal any significant differences.

5. Bad Pixels

Bad pixels are usually defined as pixels with less than 10% of the response of average pixels or hot pixels with dark currents that make them nonlinear after 1000 seconds of integration. For the NHUDF this definition was expanded to include a more stringent limit on the allowed dark current. The more stringent definition increased the number of bad pixels from 170 to 1012. This still, however, left warm pixels with significant dark current. This dark current varied with detector temperature and had significant poisson variation compared to the galaxy flux. After dark image subtraction these pixels often produced either spikes or holes in the images, depending on whether their temperature change and poisson error made them higher or lower than the median represented by the dark image. These pixels usually appeared as singular high contrast dark or bright spots.

These pixels are identified by comparing their contrast to the surrounding pixels with the maximum contrast in the well known NICMOS Point Spread Function (PSF). In a single image a pixel value was declared bad if all three of the following criteria were met.

1. The value of the pixel is greater than the 3σ noise value

2. Its contrast relative to the 8 surrounding pixels is higher than the maximum contrast from a centered point source.

3. The average value of the surrounding pixels was less than a preset number of σ.

The last criterion was established so that the centers of unresolved galaxies were not corrected if poisson noise increased their contrast above the maximum contrast of the PSF.

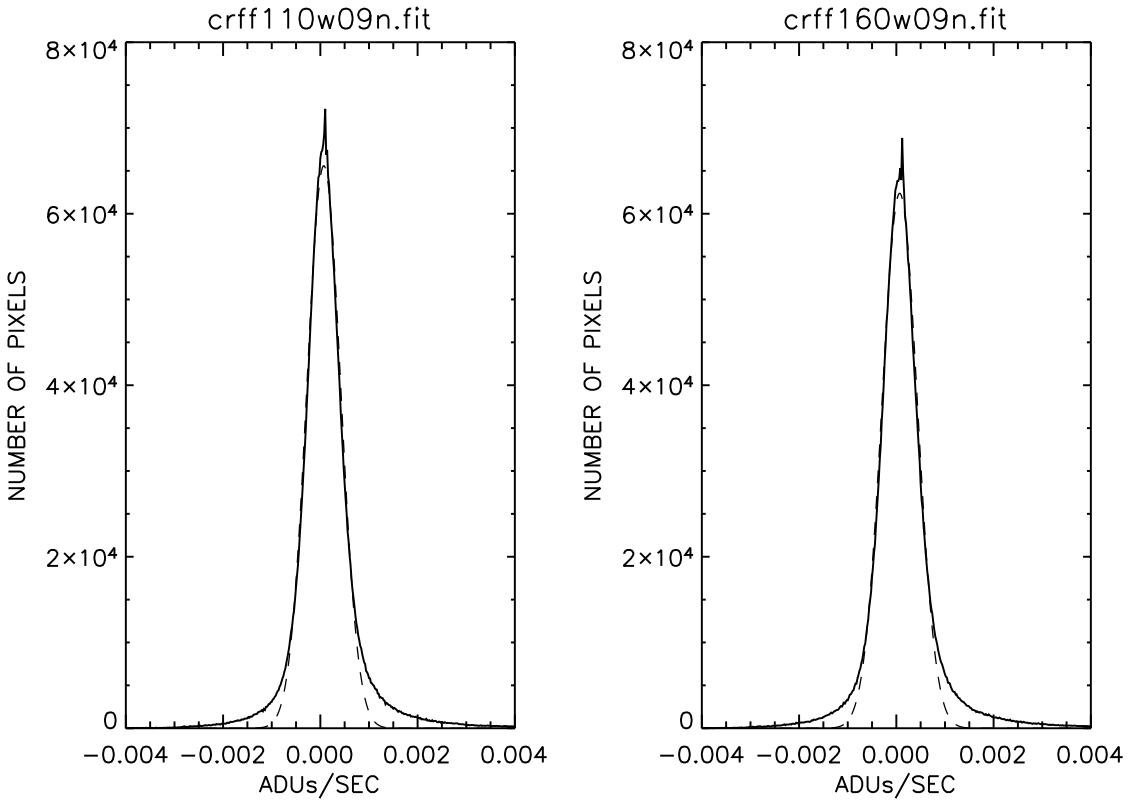

Figure 1: This is the distribution of pixel values for the F110W and F160W images in ADUs/sec. along with a Gaussian fit to the distribution shown by the dashed line. The long positive tail on each distribution is the contribution from real sources.

6. Quadrant Bias Removal

The NICMOS detectors are divided into 4 quadrants of 128x128 pixels, each with an independent readout amplifier. This was done to provide redundancy against amplifier failure. This sometimes results in a phenomenon called quadrant bias which is a DC offset of the voltages of all pixels in a detector quadrant. If there is an offset, the flat field correction imprints its pattern into the DC bias. The flat field correction function has significant power on small spatial scales which can be detected by a procedure developed by Mark Dickinson and described in the HST NICMOS data handbook (Mobasher et al. 2004). The procedure adds and subtracts a range of DC offsets to the image before the flat field procedure. The fluctuations in the image are measured after the flat field is applied and the offset that produces the minimum fluctuation is used to correct the quadrant bias. The fluctuation is measured by the FWHM of the distribution of all pixel values. Figures 1 and 2 show the distribution of pixel values and the fluctuation distribution measured by a Gaussian fit to the pixel value distribution for a typical image in the NHUDF.

This procedure works well for images that contain faint objects and where the majority of pixels are background. Figure 1 shows that this is the case for the NHUDF. For images with bright objects the quadrant bias offsets may not be significant.

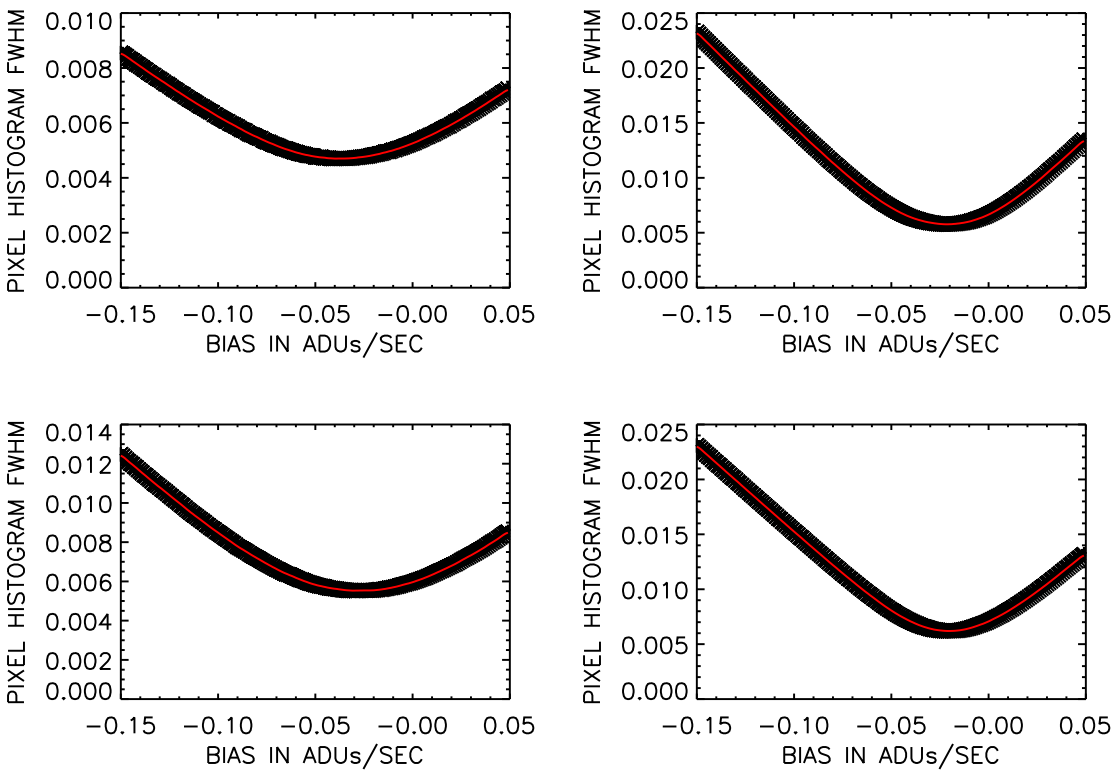

Figure 2: This shows the value of the fluctuations as measured by the FWHM of the pixel value distribution versus the offset applied in the quadrant bias removal procedure.

7. Background Subtraction

As discussed in Section 3 the sky background flux is higher than the fluxes from the majority of objects and much higher than the 1σ fluxes. Careful background subtraction is therefore very important for detecting objects and for accurate photometry. The primary source of the background is zodiacal light which is primarily reflected light rather than thermal emission at the NICMOS wavelengths. The background is determined by taking a median of all of the images in a filter. The 3x3 tiling and the dithering of images at each tile position ensures that the median measures the smooth background rather than the average of the source flux. The median produced by this procedure is very flat and shows no trace of source structure to the noise limit of the image. Since the zodiacal flux is time and position variable, the median background images were produced for each of the two epochs of observation. The background image was subtracted from each image before the images were combined in the drizzle process.

8. SAA Persistence

In general SAA persistence is an important problem for the observation of faint objects. Correction methods for images with persistence have been developed by STScI and are described in the NICMOS handbook. The NHUDF, however, benefited from an excellent scheduling effort by Beth Perriello and did not have any SAA impacted orbits. Occasional large cosmic ray hits did occur and were generally handled by masking out the affected areas of the image as discussed in sections 9. and 10..

9. Visual Inspection

Before the images were drizzled all 288 (144 F110W and 144 F160W) were visually inspected over a range of stretches for any artifacts that might affect the quality of the final image. The most common artifacts found in the images were persistent images from cosmic rays and satellite passages through the field of view. The cosmic ray events appear to be showers from high energy protons impacting the instrument and producing an intense shower of byproducts that affect a large area of the detector. Less than 10 images per filter were affected by this phenomenon. This is a very important step in the analysis due to the order of the integrations during an orbit. Each orbit contains a F110W image followed by a F160W image. A cosmic ray hit before the start of integrations in an orbit will produce a persistent image that will be strongest in the F110W image and will have faded to produce a weaker F160W image. The result will be an optical dropout source with a blue F110W to F160W color, the signature of a high redshift source. Every optical dropout candidate was inspected to make sure that the resultant flux did not come from such a single event. This resulted in the elimination of one source that appears in the Treasury Version 1 images but is removed from the Version 2 images. In addition any obvious noise sources were also identified.

10. Masking

All artifacts identified by visual inspection or other means were masked out in the images where they occurred by setting the flux to zero and setting the weight to zero in the drizzle process. In addition to the individual image masks the bottom 15 rows of all of the images were masked out. This region suffers slight vignetting which is of no consequence in most images but produces a region of significantly higher noise in these images. A small triangular region in the upper right hand corner of all images was also masked out. This is

a region of strongly varying pixel response which again has negligible effect in most images but produces a higher noise region in the NHUDF images. The common mask and the individual image masks are part of the NHUDF Version 2 data set available from MAST. It is highly recommended that any use of the NHUDF images for statistical purposes should take into account the varying weight across the images. This is mainly due to the variance in detector sensitivity but is also affected by the masked area.

11. Photometric Accuracy

Photometric accuracy of the images is extremely important, particularly for those studies that match template Spectral Energy Distributions (SEDs) to the combination of ACS and NICMOS measurements in the HUDF. There are also discussions of the photometric accuracy of NICMOS in other contributions to these proceedings (see the contributions by deJong and Mobasher).

11.1. Relative Accuracy

The relative accuracy of the NHUDF images is defined as the relative accuracy between the fluxes of objects in an image and the accuracy of the F110W to F160W color. The dithering and tiling has smoothed out the large QE variations in the NICMOS camera 3 detector, however, that is already well handled by the flat field correction. More importantly the dithering has mitigated the effect of intra pixel QE variations. This results in a relative accuracy of better than 5% with some areas of high signal to noise having significantly better accuracy.

11.2. Absolute Accuracy

Presentations and discussions at the workshop raised questions about the absolute calibration of the camera 3 F110W and F160W filters. As a result of this discussion the author revisited the initial calibration, particularly in light of the much more extensive standard star observations that are available now, which were not available at the time of the initial calibration. This is material that was not presented at the workshop but is an important contribution to NICMOS calibration and is therefore included in the proceedings. As a result of the work below the new calibration of the NICMOS camera 3 F110W and F160W filters gives the sensitivity of the F110W band as 1.70×10^{-6} Janskys per ADU per second and the F160W band as 2.06×10^{-6} Janskys per ADU per second. This means that a star with a Vega spectrum will have that flux in Janskys at the effective wavelength of the filters which are 1.103 and 1.595 microns respectively for the F110W and F160W filters.

11.3. New Photometric Analysis

The recalibration utilizes the observations of the solar analog star P330-E since the installation of the NCS in February of 2003. The star is assumed to have the same spectrum as the solar spectrum determined by Thuillier et al. (2003). The calibration also assumes that the photometric data on Vega by Campins, Rieke and Lebofsky (1985) and Bessell (1990) are correct.

P330-E is one of the solar analog stars observed by Persson et al. (1998) as part of the calibration program for NICMOS. The first step in the calibration process is the calculation of the effective wavelength of the Persson and NICMOS filters. The effective wavelength is defined by

$$\lambda_e = \frac{\int \lambda T(\lambda) d\lambda}{\int T(\lambda) d\lambda} \tag{1}$$

where $T(\lambda)$ is the transmission of the filter only and does not include the quantum efficiency of the detector. The effective wavelengths of the Persson filters were calculated from the transmission functions given in Persson et al. (1998) and the NICMOS effective wavelengths from the measured transmission of the filters before flight. Table 1 gives the measured values of the effective wavelengths.

Table 1: The effective wavelengths of the filters in microns.

Filter	Persson J	Persson H	Persson K	NICMOS F110W	NICMOS F160W
Effective λ	1.24	1.554	2.257	1.103	1.595

The quantum efficiency of the Persson observation detector was measured by hand from Figure 4 of Persson et al. (1998) and multiplied by the filter transmission to produce the response functions of the Persson system.

The next step was the transfer of the Campins et al. (1998) and Bessell (1990) photometry to the Persson system by fitting the Campins and Bessell photometry with a cubic spline fit and calculating the expected signal using the Persson response function. The Persson 0 magnitude values were found to be 1606, 1019 and 644 Janskys for J, H and K taking into account the 0.02 magnitude difference between the Arizona and Persson systems.

Using the Persson measurement of P330-E the flux at the effective wavelengths was determined. These fluxes, along with the scaled fluxes of the Thuillier et al. (2003) solar spectrum are listed in Table 2. The Thuillier et al. (2003) solar spectrum has some differences from the previously used solar analog spectrum as is shown in Figure 3.

Table 2: P330-E photometry and the scaled solar flux to match it.

Filter	P330-E (Jy)	Solar (Jy)
J	3.077×10^{-2}	3.138×10^{-2}
H	2.53×10^{-2}	2.453×10^{-2}
K	1.749×10^{-2}	1.756×10^{-2}

The scaled solar spectrum was then integrated over the F110W and F160W response function to determine the expected number of ADUs per second from the star. These values were then compared to the P330-E observations from proposal IDs 9325, 9639, 9995, and 10381 which together contain 182 observations of the star in the F110W and F160W filters with camera 3. To eliminate background flux a median of all 182 images was taken and then subtracted from the individual images. The median was inspected for any residual source flux and none was found since the location of the image was essentially randomly placed between the images.

The median subtracted images were then analyzed using SExtractor (SE) with 5 different photometric apertures of 2.7, 5, 8, 15 and 30 pixels. The pixels are 0.02 arc seconds in size. The smallest aperture was chosen to match the smallest aperture in the Version 2 Treasury data for the NHUDF. Initially an 80 pixel aperture was also used but its flux was identical within the noise to the 30 pixel aperture indicating that all of the flux was included in the 30 pixel aperture. Table 3 gives the aperture functions derived from these measurements and Figure 4 shows the functions.

Based on these measurements the sensitivity of the F110W filter was reduced by a factor of 0.925 and the sensitivity of the F160W filter by 0.910. The F110W and F160W fluxes in the Version 2 Treasury NHUDF should be increased by the inverse of these reductions. It should be noted that the aperture fluxes listed in the catalog are strictly the flux observed

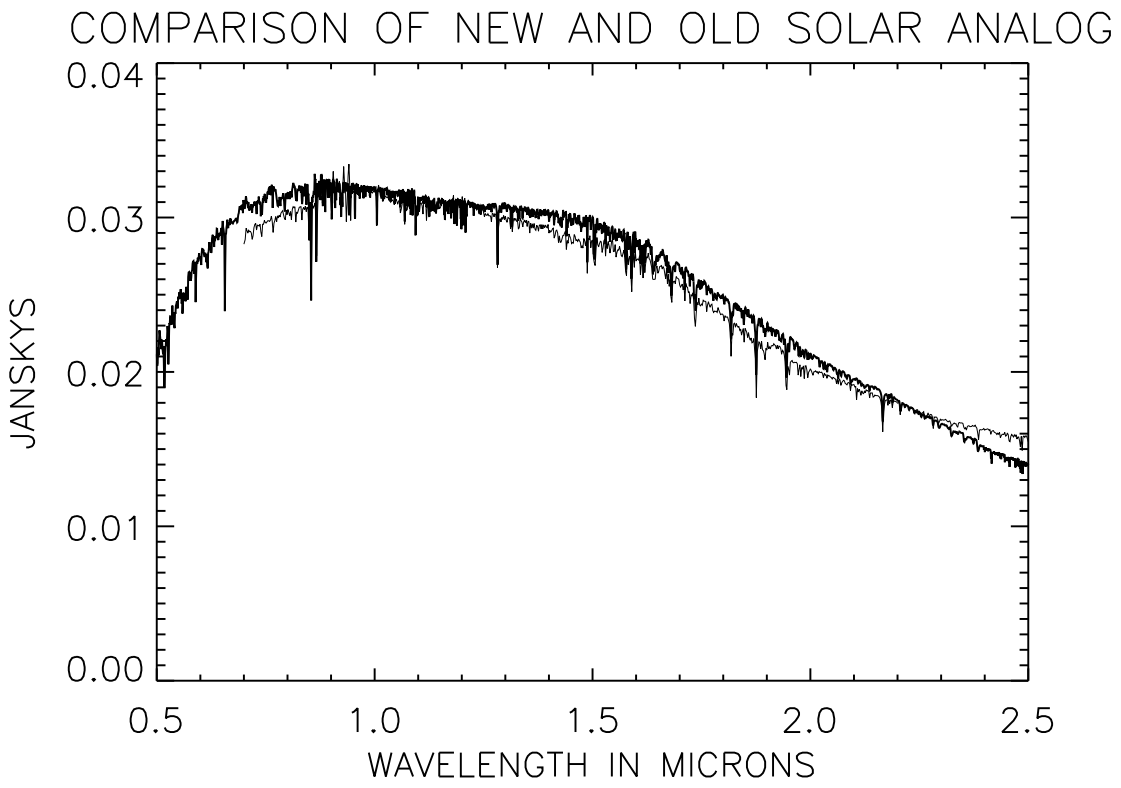

Figure 3: Comparison between the solar spectrum from Thuillier et al. 2003 (thick line) and the old solar analog spectrum (thin line).

Table 3: The aperture functions for the NICMOS Camera 2 F110W and F160W filter for a solar spectrum star. The aperture sizes are in arc seconds.

Filter/Ap.	0.54	1.0	1.6	3.0	6.0
F110W	0.700	0.889	0.950	0.990	1.0
F160W	0.645	0.856	0.924	0.978	1.0

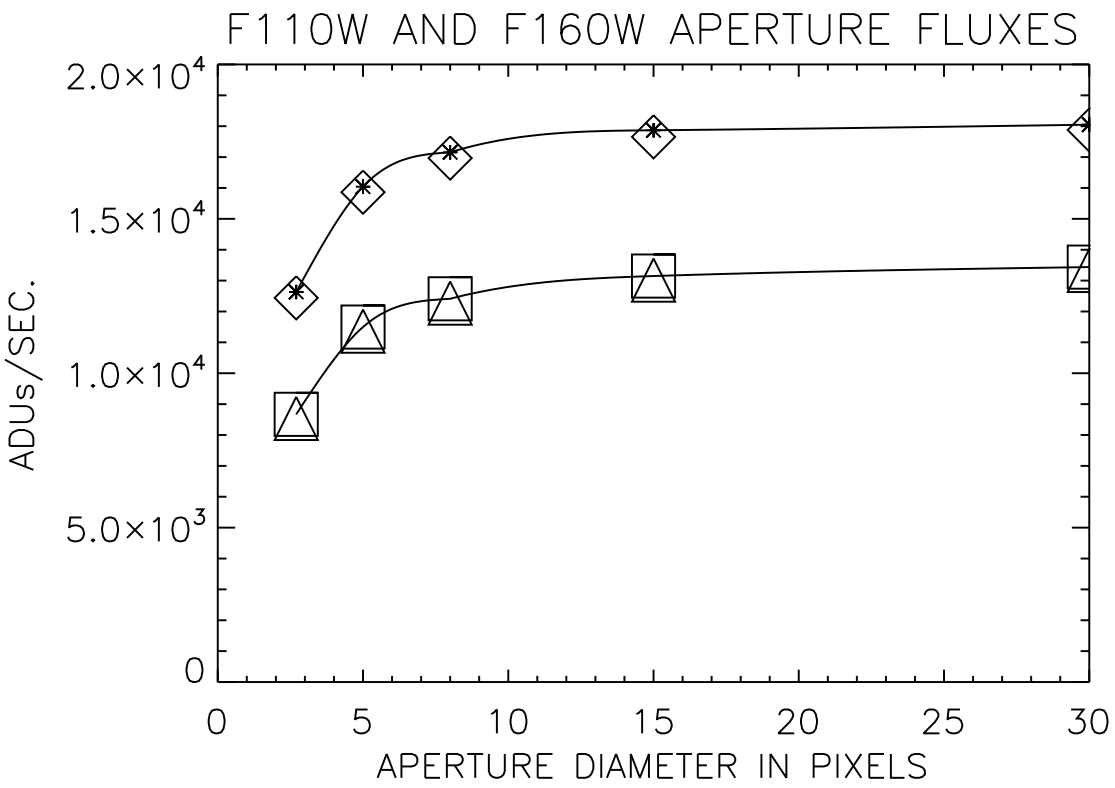

Figure 4: The F110W and F160W aperture functions. The diamonds and asterisks are the average and median F110W fluxes. The triangles and squares are the F160W average and median fluxes.

in those apertures and they should be increased by the inverse of the aperture function to account for the full flux when comparing the flux in the apertures relative to the ACS flux. Aperture corrections for the ACS observations are given in Sirianni et al. (2005).

12. Possible Non-Linearity

The possibility of a non-linear response in NICMOS observations has been raised in this workshop. Refer again to the contributions by deJong and Mobasher. At the time of this submission the author has not been able to independently check the validity of this effect but users of the catalog should be aware of the possibility of its existence. After this check is completed the author expects to submit a version 3 of the NHUDF observations to MAST taking into account all of the calibration issues discussed here.

Acknowledgments. This paper is based on data from observations with the NASA/ESA Hubble Space Telescope, obtained at STScI which operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555.

References

Bessell, M. S. 1990, PASP, 102, 1181

Campins, H., Rieke, G. H. & Lebofsky, M. J. 1985, AJ, 90, 896

Mobasher et al. 2004, *HST Data Handbook for NICMOS*, version 6.0, (Baltimore: STScI), http://www.stsci.edu/hst/nicmos/documents/handbooks/DataHandbookv6

Persson, S. E. et al. 1998, AJ, 116, 2475

Sirianni, M. et al. 2005, PASP 117, 1049

Thompson, R. I. et al. 2005, AJ, 130, 1

Thuillier, G. et al. 2003, Solar Physics, 214, 1

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

A Critical Test of NICMOS Photometry

Bahram Mobasher, Roelof de Jong, Adam Riess, L. Eddie Bergeron & Keith Noll

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. We compared photometry for the *same* sources as measured by NICMOS and ground-based ISAAC detectors over a large dynamic range. Over a range of magnitudes from 8 - 25, we find a linear offset in the F110W filter with respect to ground-based J-band once the data are corrected for differences in J-band filters. No offset or non-linearity is observed between F160W and ground-based H-band data.

Comparing NICMOS photometry for faint objects in HUDF, reduced in two independent ways, we find a difference of 0.04-0.08 mag between the two datasets. This is completely independent from the filter shapes and detector sensitivities and only depends on the way the cosmic ray and sky subtraction is applied. Allowing for dependence of flat fields on zero-points, slightly increases the offset.

1. Introduction

A through examination of the photometric accuracy of the Near Infrared Camera and Multi-Object Spectrometer (NICMOS) over a large dynamic range is needed to test dependence of the photometric sensitivity to flatfield and dark estimates, the zero-points, sky subtraction and different reduction techniques. In this paper we carry out this investigation in two different ways. Firstly, we compare photometric observations of the same sources using NICMOS and another, independent, near-IR detector over a similar wavelength range. To ensure high dynamic range, we utilize both bright calibration stars and faint compact galaxies. This allows an examination of the non-linearity of NICMOS and a test of its zero-point. Secondly, we perform a comparison of the same NICMOS dataset, reduced through two different techniques. This provides a comprehensive test to examine the sensitivity of NICMOS photometry to different flatfield and dark estimates and background subtraction techniques. Given the importance of photometric accuracy at low flux levels, we investigate this at the very faint level.

2. A Test of NICMOS Non-linearity over Large Dynamic Range

To carry out the non-linearity test, we compare NICMOS photometry in F110W and F160W filters with photometry for the same objects, taken with ground-based near-IR detectors. To ensure high dynamic range, we use both bright calibration stars and faint galaxies. We took the bright stars from Stephens et al (2000), for which both NICMOS (F110W and F160W) and ground-based (JH) data are available. For the faint galaxies, we use NICMOS (NIC3) images in the Hubble Ultra Deep Field (HUDF), for which ground-based J and H-band data exist from the Infrared Spectrometer and Array Camera (ISAAC) on the Very Large Telescope (VLT).

We took steps to minimize systematic effects between the NICMOS and ground-based data. In the high-flux regime, we applied corrections from Stephens et al (2000) to convert magnitudes to the same system. In the low-flux regime, the filter response functions for the F160W and H-band are similar, resulting in negligible difference between the magnitudes.

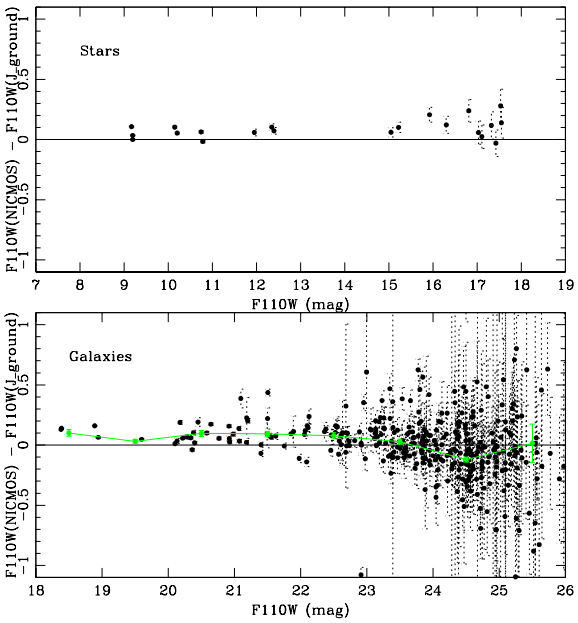

Figure 1: Comparison between NICMOS (F110W) and ground-based J-band photometry for stars (top panel) and galaxies (bottom panel). Stars are taken from Stephens et al (2000) and galaxies from HUDF with available NICMOS and ground-based data. Flat horizontal line at 0 represents agreement between magnitudes for point source (stars) and extended (galaxies) objects. The green line represents the data binned in magnitude intervals. All the magnitudes are corrected for differences between the filter response functions.

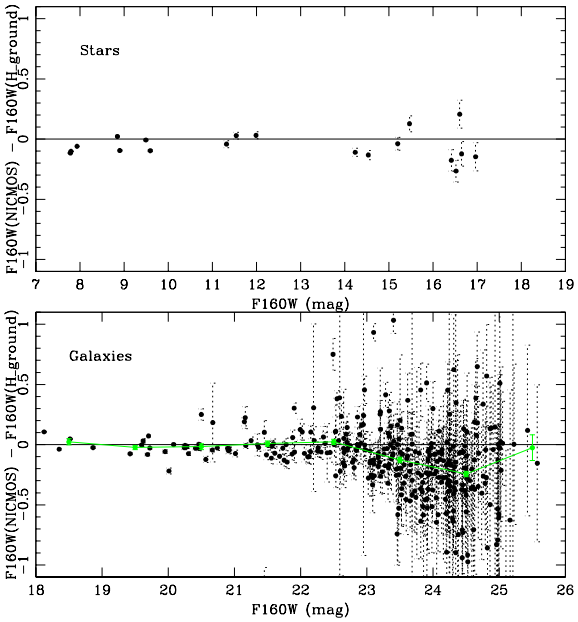

Figure 2: The same as Figure 1 for F160W filter.

However, the F110W filter extends to shorter wavelengths than the standard ISAAC J-band, making a direct comparison more difficult. We corrected for this using known photometric redshifts and spectral types for HUDF galaxies and estimating corrections to magnitudes, ranging from 0.2 to 0.6 magnitude, using their best-fit Spectral Energy Distributions.

We compare photometry for F110W *vs.* J (Figure 1) and F160W *vs.* H (Figure 2) over a dynamic range from 8 to 25 magnitude. The faintest six magnitudes over this range are mostly dominated by background counts. In the case of F160W, the magnitudes from the two detectors agree quite well, demonstrating good consistency for F160W photometry. Moreover, the continuity between the magnitudes for stars and galaxies is better than 0.1 magnitude. We find the NICMOS F110W magnitudes are consistent with ground-based J-band data in the low-flux regime, after reducing dispersion in magnitudes by averaging a large number of objects, divided into magnitude intervals. This comparison depends on the accuracy of photometric transformations, corrections due to differences in filter response functions and matching the apertures between the ground-based and NICMOS observations. In the high flux regime, the F110W magnitudes are 5-10% fainter than ground-based J-band magnitudes.

The main cause of the excess faintness for NICMOS magnitudes, also found by Bohlin, Linder & Riess (2005), is due to a non-linearity in NICMOS photometry, as is recently confirmed by de Jong et al. Using the recently obtained NICMOS data, the size of the non-linearity is now calibrated.

3. NICMOS Photometric Comparison From Different Reduction Techniques

There are two versions of the NICMOS HUDF data, reduced completely independently by Thompson et al (2005) and Robberto et al (2006). The images are reduced following standard procedure, using library flat fields and dark images. The main difference between the two reduction procedures is the cosmic ray removal and sky subtraction techniques. Robberto et al (2006) uses an iterative technique to perform cosmic ray removal and sky subtraction while Thompson et al (2005) generate the median filter image from all the HUDF NICMOS observations and use that as the background image. In order to compare the resulting photometric data from the two reduced images, we perform photometry on both the images in *exactly* the same way.

We used SExtractor to detect sources from Robberto et al (2006). We then performed aperture photometry on the position of the detected sources on both Thompson et al/ and Robberto et al/ images. We used the same aperture size, local background subtraction procedure and zero-points. We compare the results from the two analysis in Figures 3 (J-band) and 4 (H-band). In this analysis the difference between the two measurements is entirely due to the reduction techniques used by the two teams. We find a difference of 0.04-0.08 mag in the photometry of faint galaxies, resulting from differences in the reduction techniques. One source of uncertainty here is that due to dependence of zero-points on flat fields, one needs to use consistent zero-points. This might be the case as Thompson's reduction used old zero-points, calculated before installation of the NICMOS Cooling System. Using consistent photometric zero-points slightly increases the offset between J (Figure 3- bottom panel) and H (Figure 4- bottom panel) band magnitudes.

4. Conclusions

We compared photometry for the *same* sources as measured by NICMOS and ground-based ISAAC detectors over a large dynamic range. Over a range of 8-25 magnitudes, we find a linear offset in the F110W filter with respect to ground-based J-band once the data are

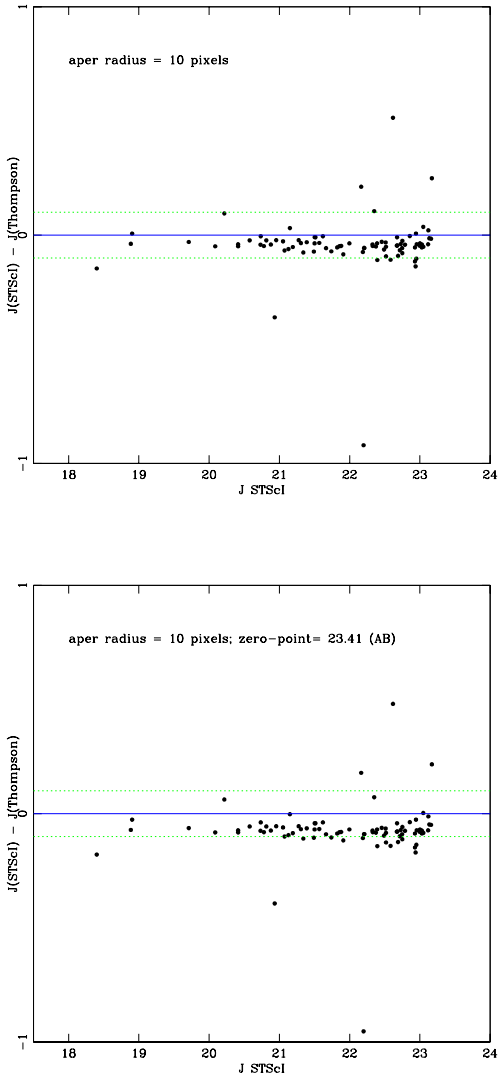

Figure 3: Comparison between two independent reductions of the HUDF NICMOS data from Robberto et al (2006)-(STScI) and Thompson et al (2005)- (Thompson). The magnitudes are measured over 10 pixel diameter with only the brighter objects used. Top Panel: the same photometric zero-points are used. Bottom Panel: photometric zero-points were used consistent with the data

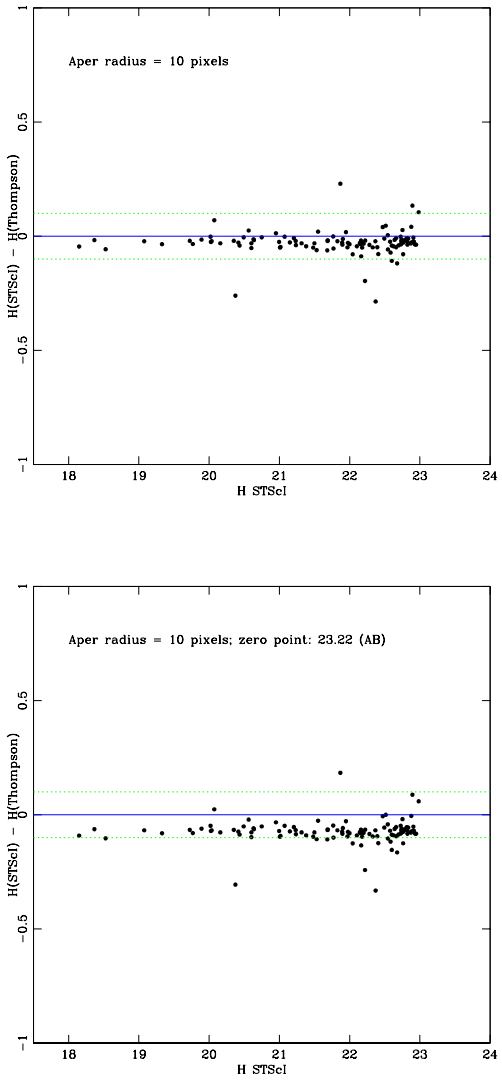

Figure 4: The same as Figure 3, for the F160W band.

corrected for differences in J-band filters. No offset or non-linearity is observed between F160W and ground-based H-band data.

Comparing NICMOS photometry for faint objects in HUDF, reduced in two independent ways, we find a difference of 0.04-0.08 mag between the two datasets. This is completely independent from the filter shapes and detector sensitivities and only depends on the way the cosmic ray and sky subtraction is applied. Allowing for dependence of flat fields on zero-points slightly increases the offset.

References

Bohlin, R., Lindler, D. J. & Riess, A. 2005, *Instrument Science Report* NICMOS 2005-00 (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Robberto et al. 2006, in preparation

Stephens, A. W., Frogel, J. A., Ortolani, S., Davies, R., Joblonea, P. Renzini, A., Rich, R. 2000 ApJ 119, 419

Thompson, R. et al Astron. J. 2005

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Polarimetry with NICMOS

Dean C. Hines[1], Glenn Schneider[2]

Abstract. NICMOS cameras 1 and 2 (with ~ 43 and ~ 76 mas/pixel, respectively) each carry a set of three polarizing elements to provide high sensitivity observations of linearly polarized light. The polarizers are bandpass limited and provide diffraction-limited imaging in camera 1 at 0.8 - 1.3μm, and in camera 2 at 1.9-2.1μm.

The NICMOS design specified the intra-camera primary axis angles of the polarizers to be differentially offset by 120°, and with identical polarizing efficiency and transmittance. While this ideal concept was not strictly achieved, accurate polarimetry in both cameras, over their full ($11''$ and $\sim 19.2''$ square) fields of view was enabled through ground and on-orbit calibration of the as-built and *HST*-integrated systems.

The Cycle 7 & 7N calibration program enabled and demonstrated excellent imaging polarimetric performance with uncertainties in measured polarization fractions $\leq 1\%$. After the installation of the NICMOS Cooling System (NCS), the polarimetric calibration was re-established in Cycle 11, resulting in systemic performance comparable to (or better than) Cycle 7 & 7N.

The NCS era NICMOS performance inspired the development of an earlier conceived, but non-implemented, observing mode combining high contrast coronagraphic imaging and polarimetry in camera 2. This mode was functionally tested in the Cycle 7 GTO program, but without a detailed characterization of the instrumental polarization induced by the coronagraph, proper data calibration was not possible. To remedy this shortfall and to enable a new and powerful capability for NICMOS, we successfully executed a program to calibrate and commission the "Coronagraphic Polarimetry" mode in NICMOS in Cycle 13, and the mode was made available for GO use in Cycle 14.

We discuss the data reduction and calibration of direct and coronagraphic NICMOS polarimetry. Importantly, NICMOS coronagraphic polarimetry provides unique access to polarized light near bright targets over a range of spatial scales intermediate between direct polarimetry and ground-based (coronagraphic) polarimetry using adaptive optics.

1. Thermal Vacuum Tests

The NICMOS polarimetry systems were characterized on the ground during thermal vacuum tests using a light source that fully illuminated the field of view with completely linearly polarized light (i.e., p(%) = 100%) and with position angles variable in 5° increments. The polarizers in camera 1 and camera 2 are labeled by their nominal position angle orientation: POLS* for camera 1 and POLL* for camera 2. Camera 3 has three grisms that potentially act as polarizing elements, and were characterized as well. However, because the grisms

[1]Space Science Institute, 4750 Walnut Street, Suite 205 Boulder, CO 80301

[2]Steward Observatory, The University of Arizona, 933 N. Cherry Ave., Tucson, AZ 85721

reside in the NIC3 filter wheel, they cannot be used with either the NIC1 or NIC2 polarizers and are unsuitable for spectropolarimetry.

The primary results of these thermal vacuum tests include:

- Each polarizer in each camera has a unique polarizing efficiency[1], with POL120S having the lowest at $\epsilon_{\rm POL120S} = 48\%$.
- The angular offsets between the polarizers within each filter wheel differ from their nominal values of 120°.
- The instrumental polarization caused by reflections off the mirrors and optical baffles in the optical train is small ($\leq 1\%$).
- The grisms indeed act as partial linear polarizers, with the long wavelength grism (G206) producing the largest variation in intensity ($\sim$5%) for 100% linearly polarized light. If one measures the equivalent width of an emission line against a polarized continuum, the measurement will have different values depending on the orientation of the spacecraft *w.r.t.* the position angle of the polarized continuum. This effect scales with the source polarization, so should not be important for objects of low polarization (p(%) $\leq 10\%$).

2. The Algorithm for Reducing NICMOS Polarimetry Observations

The thermal vacuum results showed that the standard reduction algorithm listed in basic optics text books would not work for NICMOS data. Instead, a more general approach was required (Hines, Schmidt & Lytle 1997; Hines 1998; Hines, Schmidt & Schneider 2000; Hines 2002).

At any pixel in an image, the observed signal from a polarized source of total intensity I and linear Stokes parameters Q and U measured through the k^{th} polarizer oriented at position angle φ_k is

$$S_k = A_k I + \epsilon_k (B_k Q + C_k U) \ . \tag{1}$$

Here,

$$A_k = \frac{t_k}{2}(1 + l_k), \quad B_k = A_k \cos 2\varphi_k, \quad C_k = A_k \sin 2\varphi_k \ , \tag{2}$$

ϵ_k is the polarizing efficiency, t_k is the fraction of light transmitted through the polarizer for a 100% linearly polarized input aligned with the polarizer axis, and l_k is the "leak" – the fraction of light transmitted through the polarizer (exclusive of that involved in t_k) when the incident beam is polarized perpendicular to the axis of the polarizer. These quantities are related under the above definitions, $\epsilon_k = (1 - l_k)/(1 + l_k)$.

This treatment can be shown to be equivalent to other approaches, once appropriate transformations are made (Mazzuca, Sparks, & Axon 1998; see also Sparks & Axon 1999).

The values of t_k were determined initially by the filter manufacturer from witness samples, and were not accurately remeasured during thermal vacuum tests. However, on-orbit observations of the unpolarized and polarized standard stars enables refinement of these numbers. Adopted characteristics of the individual polarizers and algorithm coefficients derived during and applicable to Cycle 7 & 7N are listed in Table 1 of Hines, Schmidt & Schneider (2000), the Cycle 11 coefficients are listed in Hines (2002), and are also available

[1]Polarizing efficiency is defined as $\epsilon = (S_{\rm par} - S_{\rm perp})/(S_{\rm par} + S_{\rm perp})$, where $S_{\rm par}$ and $S_{\rm perp}$ are the respective measured signals for a polarizer oriented parallel and perpendicular to the position angle of a fully polarized beam.

in the NICMOS instrument manual and the *HST* Data Handbook. The primary coefficients are also listed in Table 1 in this contribution.

After solving the system of equations (eq. 1) for the Stokes parameters at each pixel (I, Q, U), the percentage polarization (p) and position angle (θ) are calculated in the standard way:

$$p = 100\% \times \frac{\sqrt{Q^2 + U^2}}{I}, \quad \theta = \frac{1}{2}\tan^{-1}\left(\frac{U}{Q}\right) \ . \tag{3}$$

Note that a 360° arctangent function is assumed.

This algorithm has been tested by the NICMOS Instrument Definition Team (IDT) and by the Space Telescope Science Institute (STScI) on several data sets. An implementation has been developed by the IDT, and integrated into a software package written in IDL. The package is available through the STScI website[2] and is described by Mazzuca & Hines (1999). An update with improved graphical user interface will be available for Cycle 15.

3. The Cycles 7, 7N & 11 Polarimetry Characterization Programs

Observations of a polarized star (CHA-DC-F7: Whittet et al. 1992) and an unpolarized (null) standard (BD + 32 3739: Schmidt, Elston & Lupie 1992) were obtained with NIC1 and NIC2 in Cycle 7 & 7N (CAL 7692, 7958: Axon PI) and in Cycle 11 (CAL 9644: Hines PI). Data were obtained at two epochs such that the differential telescope roll between observations was $\approx 135°$. This removes the degeneracy in position angle caused by the pseudo-vector nature of polarization. The second epoch observations in Cycle 7N and all of the observations in Cycle 11 used a four position, spiral-dither pattern with 20.5 pixel (NIC1) and 30.5 pixel (NIC2) offsets to improve sampling and alleviate the effects of bad pixels, some persistence, and other image artifacts. This is also the recommended observing strategy for all NICMOS, direct imaging polarimetry programs. While no dither pattern was used during the first epoch of Cycle 7, the data do not appear to suffer significantly from persistence. Observations of the highly polarized, proto-planetary nebula CRL 2688 (Egg Nebula) were also obtained to test the calibration over a large fraction of the field of view.

The observations were processed through the CALNICA pipeline at STScI using the currently available reference files. Aperture photometry was used to measure the total flux density from the stars in instrument units (DN/s) for each polarizer. The Stokes parameters were then constructed using equation (1). Since the thermal vacuum tests showed that the imaging optics themselves have little effect on the observed polarization, any measured polarization in the null standard was attributed the t_k term in the algorithm. The algorithm was applied to the polarized standard stars to check both the percentage polarization and the position angles.

The higher, yet more stable, operating temperature provided by the NCS and the three year dormancy of NICMOS may contribute to changes in the metrology of the optical system. Therefore, a program to re-characterize the polarimetry optics in Cycle 11 has been developed (NIC/CAL 9644: Hines).

The basic design of the Cycle 11 program follows the strategy undertaken in Cycle 7 & 7N, relying on observations of polarized and unpolarized standard stars as well as the proto-planetary nebula CRL 2688 (Egg Nebula). The Egg Nebula was only observed at a single epoch as a direct comparison with observations from Cycle 7 & 7N (ERO 7115: Hines; Sahai et al. 1998; Hines, Schmidt & Schneider 2000), and to evaluate any gross discrepancies across the field of view. Observations of the unpolarized standard star ($\mathrm{p}(\%)_{\mathrm{intrinsic}} = 0\%$, by definition) processed with the Cycle 7 & 7N algorithm coefficients yield $p_{\mathrm{NIC1}} = 0.7\% \pm$

[2]http://www.stsci.edu/hst/nicmos/tools/polarize_tools.html

$0.2\%, \theta_{\rm NIC1} = 74^\circ$ and $p_{\rm NIC2} = 0.7\% \pm 0.3\%, \theta_{\rm NIC2} = 73^\circ$. This suggests that the system changed. The observations of the polarized standard star is also larger ($\Delta p \approx 2\%$) in NIC1 compared with the measurements of Cycle 7 & 7N, which themselves were in excellent agreement with ground-based measurements (Hines, Schmidt & Schneider 2000).

Observations of the Egg Nebula were also analyzed with the Cycle 7 & 7N coefficients. As for the polarized standard star, the results for the Egg Nebula suggest that the polarimetry system has changed slightly[3], again by about 2% in p%.

We rederived the coefficients assuming that any measured polarization in the unpolarized (Q=U=0) standard was attributed to the t_k term in the algorithm (i.e. the variable term in the A_k coefficient). The new coefficients are listed in Table 1.

We note that we are embarking on a new study of all of the NICMOS direct imaging polarimetry calibration data in an attempt to improve the calibration to better than 0.3% (Batcheldor et al. 2006). This study was prompted by reports of a $\sim 1.5\%$ systematic residual polarization signal in some objects (e.g., Ueta et al. 2005). Indeed, Batcheldor et al. find a residual excess polarization for the ensemble of (un)polarized standard stars of p(%) $\approx 1.2\%$, which behaves as a constant "instrumental" polarization that can be subtracted (in Stokes parameters) from the observations. While insignificant for highly polarized objects, this residual will affect observations of low polarization objects such as active galactic nuclei.

Due to the non-ideal polarimetry optical system, estimation of the uncertainties in the percentage polarization require a slightly more sophisticated analysis, especially for NIC 1. Monte Carlo simulations of the uncertainties based on the NICMOS performance are given in Hines, Schmidt & Schneider (2000), and an analytical approach is given in Sparks & Axon (1999). In particular, the usual assumption that $\sigma_q = \sigma_u$ is not valid for NIC 1.

4. Coronagraphic Polarimetry with NICMOS

Coronagraphic polarimetric observations of the same unpolarized and polarized standard stars that were used in Cycle 7, 7N (CAL 7692, 7958: Axon PI) and Cycle 11 (CAL 9644: Hines PI), and observations of the face-on circumstellar disk TW Hya were each obtained at two epochs sufficiently spaced in time to permit large differential rolls of the spacecraft (i.e., field orientations *w.r.t.* the HST optics and NICMOS polarizers). At each epoch, coronagraphic polarimetric imaging was carried out at two field orientations differing by 29.9°. Following standard NICMOS Mode-2 target acquisitions, the acquired targets were observed through each of the three "Long" wavelength polarizers in Camera 2. The intra-visit repeats for the standard stars were designed to check both for repeatability (image stability) and possible image persistence (none of consequence were found). TW Hya was exposed in only two repeats for each polarizer.

To facilitate commonality in data processing, all coronagraphic imaging was done with STEP16 multiaccum sampling, though the number of samples varied between 10 to 12 to best fill the orbits for each of the targets. The data were instrumentally calibrated in an APL-based analog to the STSDAS CALNICA task using on-orbit derived calibration reference files suitable for these observations. Following the creation of count rate images, individual bad pixels were replaced by 2D weighted Gaussian interpolation (r=5 weighing radius) of good neighbors, and "horizontal striping" associated with heavily exposed targets was characterized and removed by median-collapse subtraction.

[3]The polarization structure of the Egg is not expected to change over the 5 year period between observations even though the object is known to show photometric variations.

4.1. PSF-subtracted Images

The image data for this program were obtained to investigate and demonstrate the further suppression of instrumentally scattered light via the camera 2 coronagraph (compared to non-coronagrahic polarimetry) and thus effect a reduction in the detection floor for a polarized signal in the diffuse wings of the (suppressed) scattering component of the occulted PSF core. For this purpose, the data obtained were not optimized to enable coronagraphic PSF subtraction (see, e.g., Fraquelli et al. 2004). Imaging coronagraphy under the pass band Camera 2 1.9-1.2μm polarizers will be less effective than under the 1.4-1.8μm F160W (H band) filter. Imaging coronagraphy of our calibration science target, the very-near face-on TW Hya circumstellar disk under the F160W filter was discussed in Weinberger et al. (2002). The smaller inner working angle of the coronagraph with polarizers (caused by the longer wavelength pass-band), while fully enclosing the PSF core under the polarizers, permits light from the first Airy ring to escape the coronagraphic image plane mask.

Both the unpolarized standard BD +32 -3739 and the polarized standard CHA-DC-F7 were selected to establish "truth" *w.r.t.* the instrumental polarization (and ultimately sensitivity) with polarimetric coronagraphy. They were not chosen to serve as coronagraphic reference PSF's for subtraction from TW Hya, and both calibration stars are significantly fainter than TW Hya. Additionally, the field of the unpolarized standard has a significant number of fainter stars near the target, unknown prior to the coronagraphic imaging resulting from this program. This is not a problem for the coronagraphic polarimetric calibration/validation, but as such that additionally makes it a poor choice as a PSF subtraction template star.

Nonetheless, DIRECT imaging of circumstellar dust from the Camera 2 polarizer set derived total intensity images is demonstrated (but not comprehensively in this summary) using this "template" PSF. The depth of exposure of BD+32 is insufficient to reveal the outer part of the TW Hya disk. The flux density rescaling of BD+32 is suggested as 3.48 from the 2MASS H and K magnitudes and the passband of the Camera 2 polarizers. The disk persists varying the DB+32 intensity over the full range of TW Hya:BD+32 intensities suggested in H and K, though very significant under- and over-subtraction residuals appear at the extreme of that range in the diffraction spikes, suggesting the a priori ratio was very close to "correct".

4.2. Polarimetry Comparison Between the Unpolarized Star and TW Hya

After reduction of the images in the three different polarizers, the images were processed through the NICMOS polarimetry reduction algorithm (Hines, Schmidt & Schneider 2000) to produce images of the polarization position angle (θ), the percentage polarization, and the polarized intensity ($= p \times I$). To consider the detection of polarized flux to be "well determined" we demand the polarized intensity in any pixel be at least 5-sigma above the median of the background. For this data set we choose a background region far from the target, and away from the diffraction spikes, where no contribution to the background from the target flux is seen. With this, the centrosymmetric polarization field for the near-face on TW Hya disk is obvious, and highly repeatable, in all four field orientations of the TW Hya observations.

While some residual instrumental polarization seems apparent, the lack of obvious structure in the position angle image of the unpolarized standard star is dramatically different than that of the TW Hya position angle image; the TW Hya image shows the classic "butterfly" pattern for centrosymmetric polarization cased by scattering off a face-on disk. Furthermore, the median polarization is about 2% per resolution element for the "unpolarized" star, compared with $\sim 10 - 15\%$ per resolution element for TW Hya. The results for TW Hya are shown in Figure 1. The measured polarization in the disk is consistent with measurements made from the ground with Adaptive Optical systems (D. Potter, private comm., and Hales et al. 2005).

4.3. Continued Analysis

Our analysis so far has shown that the mode is viable for future Cycles. We have sufficient data in hand to evaluate:

- The ability to remove the small instrumental polarization signature in Stokes parameter space.
- The spatial correlation of this instrumental polarization as a function of radius and azmuth with respect to the coronagraphic hole.
- The polarization sensitivity as a function of radius from the coronagraphic hole. This analysis is limited by the exposure depth of the unpolarized standard.

4.4. Recommendations

Based upon our preliminary analysis, we recommend that NICMOS coronagraphic polarimetry be released as an available mode for future Cycles. We note that the same rules that apply to direct coronagraphic imaging apply to this mode as well. In particular, the fine guidance necessary to keep the target well centered in the hole is necessary. In addition, we recommend that observations be obtained at two-roll angles to interpolate over bad pixels and develop an in situ flat field correction near the coronagraphic hole. This may be difficult to achieve in two-gyro mode.

We cannot yet tell if the current calibration data from Cycle 12 will be sufficient for GOs. At present, we would urge that GOs observe their target plus an unpolarized standard star, each at sufficient depth to obtain similar S/N in each object in each polarizer thus a minimum of two orbits for a single target program. However, a single, well exposed unpolarized standard star should be sufficient for a multi-target science program.

5. Two-Gyro Mode

Based upon the measured jitter in two-gyro mode, there should be no degradation of the direct polarimetry mode. Tests of coronagraphic imaging also suggest minimal degradation (Schneider et al. 2005, ISR2005-001). The primary impact will be the inability to obtain images at two spacecraft roll angles within a single orbit.

6. Summary

A wide variety of astronomical objects have been examined with NICMOS direct imaging polarimetry including Active Galaxies (Capetti et al. 2000; Tadhunter et al. 2000; Simpson et al. 2002), evolved stars and proto-planetary nebulae (Sahai et al. 1998; Weintraub et al. 2000; Schmidt, Hines & Swift 2002; King et al. 2002; Su et al. 2003; Ueta et al. 2005), and young stellar objects (Silber et al. 2000; Meakin, Hines & Thompson 2005). NICMOS continues to provide high fidelity, high spatial resolution imaging polarimetry. Combined with the polarimetry mode of the Advanced Camera for Surveys (ACS), *HST* provides high resolution imaging polarimetry from $\sim 0.2 - 2.1\mu$m. With the newly commissioned coronagraphic polarimetry mode, NICMOS provides further unique access to polarized light near bright targets over a range of spatial scales intermediate between direct polarimetry and ground-based (coronagraphic) polarimetry using adaptive optics.

An electronic version of the presentation given by D.C. Hines at the workshop can be found at the following URL:

http://www.stsci.edu/ts/webcasting/ppt/CalWorkshop2005/DeanHines102805.ppt

References

Capetti, A., et al. 2000, ApJ, 544, 269

Fraquelli, D. A., Schultz, A. B., Bushouse, H., Hart, H. M., & Vener, P. 2004, PASP, 116, 55

Hales, A. S., Gledhill, T. M., Barlow, M. J., & Lowe, K. T. E. 2005, ArXiv Astrophysics e-prints, arXiv:astro-ph/0511793

Hines, D. C., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 260

Hines, D.C., Schmdit, G.D. & Schneider, G. 2000, PASP, 112, 983

Hines, D. C. 1998, NICMOS and the VLT: A New Era of High Resolution Near Infrared Imaging and Spectroscopy, Pula, Sardinia, Italy, June 26-27th 1998 ESO Conference and Workshop Proceedings 55, 1998, Wolfram Freudling and Richard Hook eds., p. 63, 63

Hines, D. C., Schmidt, G. D., & Lytle, D. 1997, The 1997 *HST Calibration Workshop* with a New Generation of Instruments, p. 217, 217

King, N. L., Nota, A., Walsh, J. R., Panagia, N., Gull, T. R., Pasquali, A., Clampin, M., & Bergeron, L. E. 2002, ApJ, 581, 285

Mazzuca, L., Sparks, B., & Axon, D. 1998, *Instrument Science Report* NICMOS 98-017 (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Mazzuca, L. & Hines, D.C. 1999, *Instrument Science Report* NICMOS 99-004 (Baltimore: STScI)

Meakin, C. A., Hines, D. C., & Thompson, R. I. 2005, ApJ, 634, 1146

Sahai, R., Hines, D. C., Kastner, J. H., Weintraub, D. A., Trauger, J. T., Rieke, M. J., Thompson, R. I., & Schneider, G. 1998, ApJ, 492, L163

Schmidt, G. D., Elston, R., & Lupie, O. L. 1992, AJ, 104, 1563

Schmidt, G. D., Hines, D. C., & Swift, S. 2002, ApJ, 576, 429

Silber, J., Gledhill, T., Duchêne, G., & Ménard, F. 2000, ApJ, 536, L89

Simpson, J. P., Colgan, S. W. J., Erickson, E. F., Hines, D. C., Schultz, A. S. B., & Trammell, S. R. 2002, ApJ, 574, 95

Sparks, W.B., & Axon, D.J. 1999, PASP, 111, 1298

Su, K. Y. L., Hrivnak, B. J., Kwok, S., & Sahai, R. 2003, AJ, 126, 848

Ueta, T., Murakawa, K., & Meixner, M. 2005, AJ, 129, 1625

Tadhunter, C. N., et al. 2000, MNRAS, 313, L52

Thompson, R. I., Rieke, M., Schneider, G., Hines, D. C., & Corbin, M. R. 1998, ApJ, 492, L95

Weinberger, A. J., Becklin, E. E., Schneider, G., Chaing, E. I., Lowrance, P. J., Silverstone, M., Zuckerman, B., Hines, D. C. & Smith, B. A., 2002, ApJ, 566, 409

Weintraub, D. A., Kastner, J. H., Hines, D. C., & Sahai, R. 2000, ApJ, 531, 401

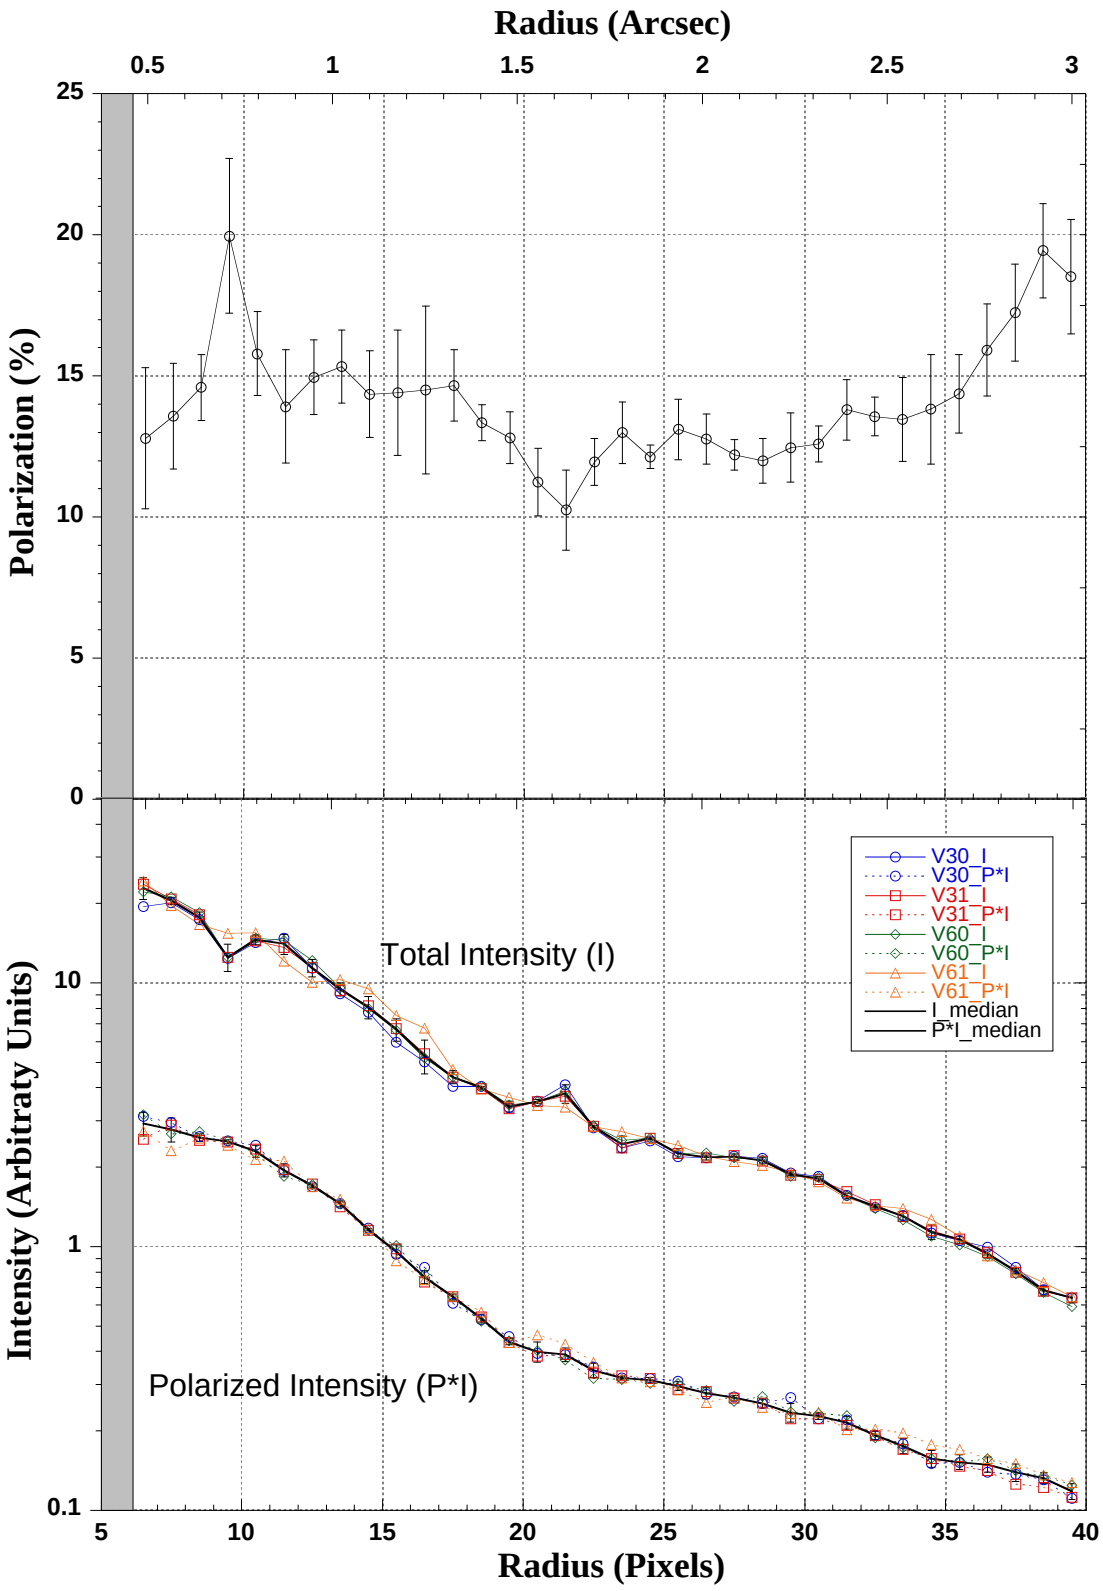

Figure 1: Coronagraphic polarimetry of TW Hya. The top panel shows the percentage polarization. The bottom panel shows the azimuthally averaged total intensity and the polarized intensity.

Table 1. Characteristics of the NICMOS Polarizers & Coefficients to Eq. (1)

Filter	φ_k	ϵ_k	t_k(Pre-NCS)[a]	t_k(NCS)[a]	Comments
POL0S	1.42	0.9717	0.7760	0.7760	Ghosts
POL120S	116.30	0.4771	0.5946	0.5934	Weak ghosts
POL240S	258.72	0.7682	0.7169	0.7173	Ghosts
POL0L	8.84	0.7313	0.8981	0.8779	...
POL120L	131.42	0.6288	0.8551	0.8379	...
POL240L	248.18	0.8738	0.9667	0.9667	...

[a]Derived from on-orbit observations of the unpolarized standard star BD+32°3739 (Schmidt, Elston & Lupie 1992).

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Differential NICMOS Spectrophotometry at High S/N

Ronald L. Gilliland[1]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. Transiting extrasolar planets present an opportunity for probing atmospheric conditions and constituents by taking advantage of different apparent radii, hence transit depth as a function of wavelength. Strong near-IR bands should support detection of water vapor via G141 spectroscopy of the bright star HD 209458 (H=6.13) by comparing in- and out-of-transit ratios of in- and out-of-band spectral intensity ratios. The reduction and analysis of science observations in which the goal is to support 1 part in 10,000, or better, development of spectral diagnostics using NICMOS grism-based spectroscopy is discussed.

1. Introduction

The study of extrasolar planetary atmospheres was initiated five years ago with detection of Sodium in the atmosphere of HD 209458b (Charbonneau et al. 2002) using STIS spectroscopy comparing spectra taken in- versus outside-of-transit. The STIS observations required use of the instrument well beyond its design goals in terms of signal-to-noise and stability, since the signal detected was only at the $2.32 \pm 0.57 \times 10^{-4}$ level for the relative increase of depth of the Sodium resonance doublet at 589.3 nm during transit.

During a planet transit light is blocked both by the solid body, and possibly by an extended atmosphere. If the extended atmosphere is opaque at some wavelengths, but transparent in others, then relatively more light (hence a deeper transit) will be blocked in the opaque wavelengths. Since atmospheres are generally very thin compared to the planet radius, this type of differential spectroscopy requires the search for small signals. Following proof of concept with the STIS results, it would now be most interesting to detect evidence for primary molecules such as water in extrasolar planet atmospheres. Although some water features exist in the optical, these are in the low QE/high fringing domain for CCDs, but much stronger bands are readily available in the near IR. Water has a strong band near 1.45 μ (see, e.g., Brown 2001 for predictions of signal strength over wavelengths of 0.3 to 2.5 μ for transmission spectroscopy of extrasolar planet atmospheres).

As part of qualifying NICMOS and its IR detector for high S/N applications, we had an earlier calibration result in which similar observations were obtained, but only during stable out-of-transit times. The results of this were released in NICMOS ISR 2003-01 (Gilliland and Arribas 2003), and concluded that differential spectrophotometry could be obtained that would support the one part in ten thousand goal. This current report provides an update based upon preliminary analysis of data from GO-9832, PI Tim Brown.

In order to obtain a detection of water vapor we initially considered how this might best be approached with NICMOS. Two primary possibilities for differential study of a feature near 1.45 μ exist: (1) There is a "Water" filter for NIC 1 with a central wavelength of 1.45

[1]Based on observations with the NASA/ESA *Hubble Space Telescope*, obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555.

μ that could be used to observe this band. (2) The NIC 3 G141 grism covering 1.1 to 1.9 μ could be used to simultaneously span wavelengths both inside and outside the water feature. Once it is realized that the signal sought may be as small as 10^{-5}, requiring detection of 10^{10} photons just to obtain sufficiently small Poisson fluctuations, *and control of all systematic errors to this same small level*, the inadequacy of using a filter becomes clear. To use the 1.45 μ filter to obtain knowledge of how this feature changes during transit, in comparison to a control band, would require cycling between filters; this in conjunction with the need to take very short exposures to avoid saturation on this H = 6.13 star would not support reaching the desired Poisson limit. Use of the grism provides two essential advantages: (1) The light is dispersed over order 100 pixels, thus allowing a much higher rate of photon collection, and even with instrument defocus HD 209458 reaches half of full well depth in only 2 seconds. (2) By simultaneously observing in- and out-of-band, the ability to correct for systematics, such as instrument sensitive drift, is far superior.

Having decided on the grism, the next decision involved dithering strategy. We explored the possibility of trailing the telescope in the cross-dispersion direction, taking about 60 seconds to trail over 200 pixels would provide the desired exposure to half of full well depth. But, ultimately, this proved impractical, if not impossible, for the ground system to support. The signal that we will pursue here is doubly differential. By this we mean that: (a) we care only about the relative change of brightness between wavelengths in-band (1.35 - 1.55 μ) to those in nearby continuum, and (b) we care only about how the latter quantity changes between times during the 3 hour transit to nearby times outside of transit. Since we care only about relative changes, it seems best to dither not at all, if the guiding is expected to be perfectly stable, or to dither only by an amount required to internally calibrate the response to the expected small guiding errors. Were we to dither the spectrum to unique pixels in an attempt to provide more accurate results this would have the effect of mixing in additional errors in the time domain; our best bet is to maintain an essentially constant pointing and take advantage of searching only for differential changes.

2. The Observations

With a full width of about 3 hours for the HD 209458b transit, it is possible to obtain equal coverage of times inside and outside transit with four contiguous HST orbits roughly centered on the transit. STIS experience suggested that instrumental stability is much poorer during the first orbit on the new target, our NICMOS experience suggests the same and we adopted the strategy of using five HST orbits to span a planetary transit, but not expecting to utilize the first HST orbit.

Even with dispersion over 100 pixels using the grism, HD 209458 would saturate NIC 3 in less than 1 second. Since there is a significant overhead associated with new MULTIACCUMS, and the minimum exposure time in any case is 0.3 seconds, it is advantageous to use moderate defocus. We chose an offset of -0.5 mm for the PAM to provide a FWHM of about 5 – 6 pixels in the cross dispersion direction. This sacrifices spectral resolution at the level of a factor of 2 – 3, but allows "longer" exposure times of 1.99 seconds to be used in which the exposure reaches only half of full well depth. Using STEP1 MULTIACCUMS with NSAMP = 4, this defocus yields spectra with peak per pixel count levels of 70,000 e-, or about 400,000 e- after summing over cross dispersion. Even with using a spectral element to spread the image over many pixels along the dispersion direction, using significant defocus, and adopting a MULTIACCUM option with a small number of samples the duty cycle (fraction of available exposure time during which photons are actually collected) is only 18%. The latter implies, that should this general approach work, then significantly fainter stars could be observed without an initially equal loss of sensitivity, since much larger duty cycles could be used.

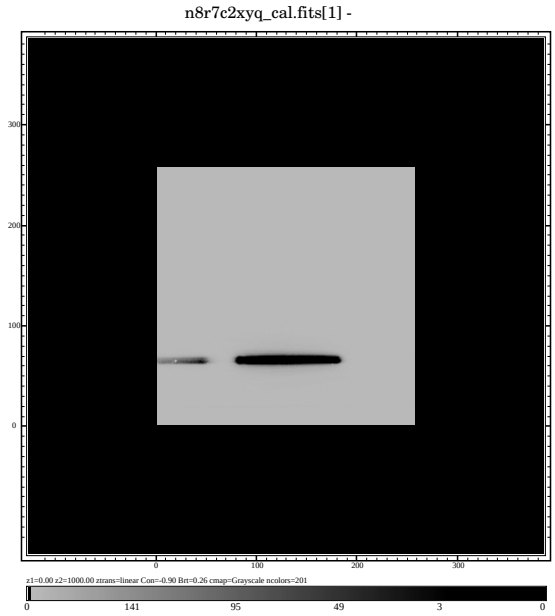

Figure 1: A G141 spectrum from GO-9832. The first order spectrum of interest is in the lower, center part of NIC 3, and was placed on a relatively clean (in terms of pixel quality) portion of the detector.

A typical example (one of 3,142 used) of a pipeline **cal.fits** data product is shown in Figure 1. The data for GO-9832 consists of 3 visits (observed in 2003, 2004, and 2005), each of 5 orbits roughly centered on the 3-hour planetary transit. About 250 exposures are obtained per HST orbit.

3. Data Reductions

The grism data are not flat-fielded as part of the pipeline reductions, since the appropriate flat would be wavelength dependent and dependent upon the distribution of sources in the field of view. In our case, with just a single point source, we have used the reference dispersion solution for G141 to provide wavelengths for each column in the first order spectrum. Then the current on-orbit flats for F108N, F113N, F164N, F166N, F187N, F190N and F196N are used to fit the wavelength dependence of the flats as a quadratic function at each pixel. The resulting quadratic was then evaluated at the appropriate wavelength by column and applied as a multiplicative correction for all pixels to be summed over in the spectral extraction. This worked very well with *rms* residuals over the 7 flats usually at the 0.001 level. Outside of the pixel area dominated by the first order spectrum, I simply adopted the F160W flat field to represent a wavelength integrated mean.

Regions well away from the spectrum (and ghosts) were chosen to allow provision of a mean, per-frame background (determined from sigma-clipped averages) which was then subtracted globally. Inspections have been made to determine if pedestal corrections need to be applied, for these ultra-high S/N observations a correction for pedestal does not seem necessary. Pixels within the spectral extraction region that are either dead or noisy have been identified and interpolated over in each individual frame (identified from both the data quality array and close inspection of the data, including 'chi-squared' maps of variance over a stack of nominally identical images).

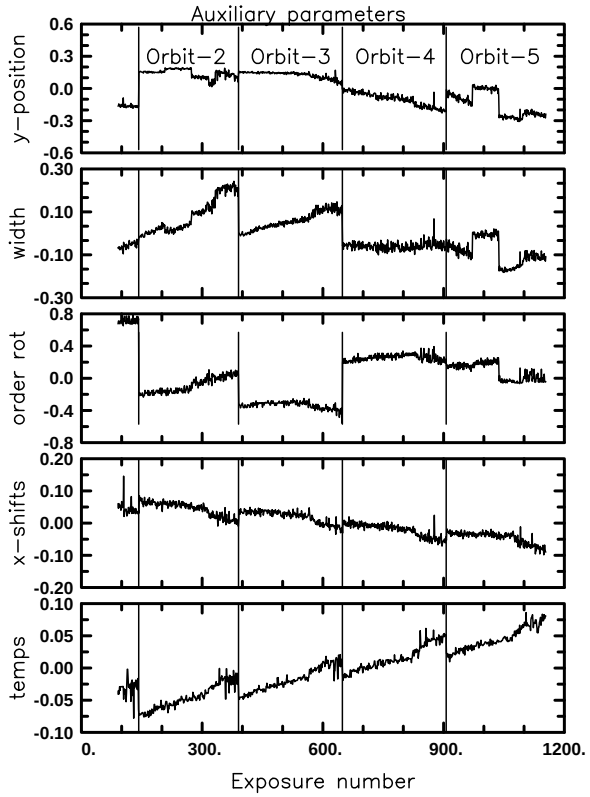

Figure 2: Variations of external parameters for from top: y-postion (cross-dispersion placement in pixels), width of the spectrum (Gaussian fit in pixels), spectrum rotation from one end to the other, x-position (from cross-correlation of extracted spectra), and temperature of the detector. Data from the 5 orbits are concatonated together for display purposes, vertical lines indicate the separate orbits.

Intensities are extracted as a simple aperture sum over the first order spectrum (up to a 20 pixel wide extraction box has been used to include most of the stellar flux). Measurements of spectral y-positions, cross-dispersion widths, and order rotation (along x) are obtained by fitting a Gaussian in cross-dispersion at many columns along the spectra, then averaging these. Exposure-to-exposure shifts in dispersion are obtained by cross-correlating individual extracted spectra with a mean spectrum. The temperature record in use was provided by Eddie Bergeron.

Two time series are of interest. The first is simply the direct intensity which shows the transit feature as a 1.6% dip. The second is the ratio of Water-band to nearby continuum band with removal of, and normalization by the mean value. Evidence of extra absorption within the water band would thus appear as a depressed value for this index during transit. Minor changes of spectrum positioning on the detector, etc., may induce changes in the time series, we therefore apply a cleaning step of decorrelating the time series variations of intensity with changes in these external parameters. Figure 2 shows the change of these over time for Visit 02. The process is: 1) evaluate a linear least-squares fit correlating the time series during orbits 2 and 5 (outside of transit) only as a function of the 5 external parameters, then 2) evaluate the model for all times and subtract from the time series in question.

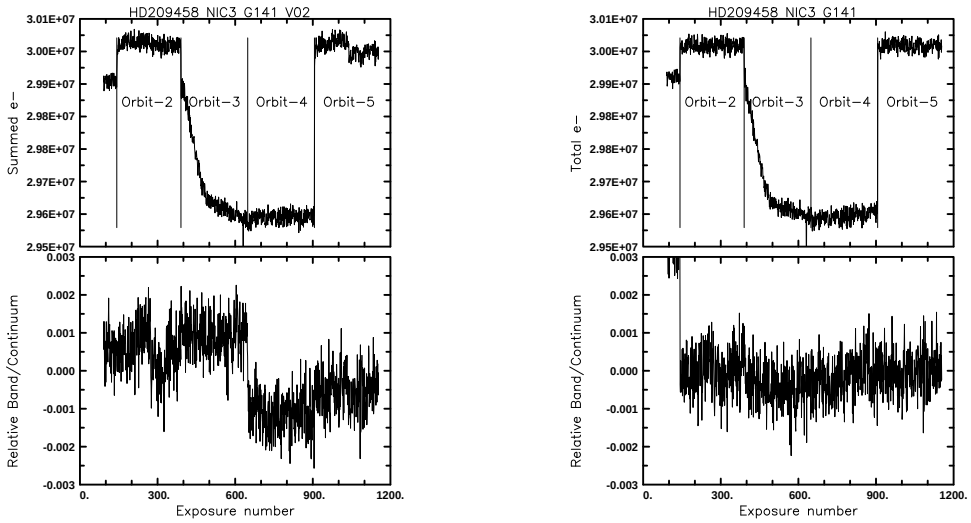

Figure 3: Left panels show the direct (top) and water-band to continuum ratio (bottom) time series after initial creation. The right panels show the same after the decorrelation step with parameters shown in Figure 2.

The before and after decorrelation time series for both the direct sum, and water-band to continuum ratio are shown in Figure 3. Except for the remaining unique behavior of orbit 1 (expected), the decorrelated time series are much cleaner.

4. Nonlinearity Test

During the planetary transit the direct intensity falls by 1.6%. Since the NICMOS detectors are known to have small nonlinearities in response to exposure level, and we wish to push these data to the levels allowed by Poisson statistics (which will be of order 10^{-5} for the ratio of in-water-band to continuum during, versus out-of-transit) an independent check of nonlinearity is advisable. An error of only $10^{-5}/0.016 = 6 \times 10^{-4}$ in the linearity correction (differential between pixels within water band and neighboring continua) would be of concern here.

During the first orbit of each visit we obtained four series of 10 exposures at each of MCAMRR, NSAMP = 6 and 7 interleaved with the standard STEP1, NSAMP = 4 observations, thus providing relative exposure times of 1.82, 2.11 and 1.99 seconds respectively. Figure 4 shows the results of combining data for all of the available test observations to constrain the linearity. After forming sums in the Water-band normalized to continua as would be done in the standard observations we find a constraint on the nonlinearity induced signal at $1.5 \times 10^{-5} \pm 1.5 \times 10^{-5}$. This result is consistent with zero and near the limiting precision available from the primary data to measure the evidence for water in the HD 209458b atmosphere. The linearity correction for count level applied in the pipeline is quite good for our purposes.

An item requiring future investigation is whether the recently detected nonlinearity with count *rate* ("Bohlin" effect) could be a limiting systematic to contend with.

5. Time Series Photometry – Comparison with STIS

Figure 5 shows overplotted results for both STIS (from Brown et al. 2001) and NICMOS time series. The STIS spectra centered at 600 nm provided a time-series *rms* of 1.1×10^{-4}

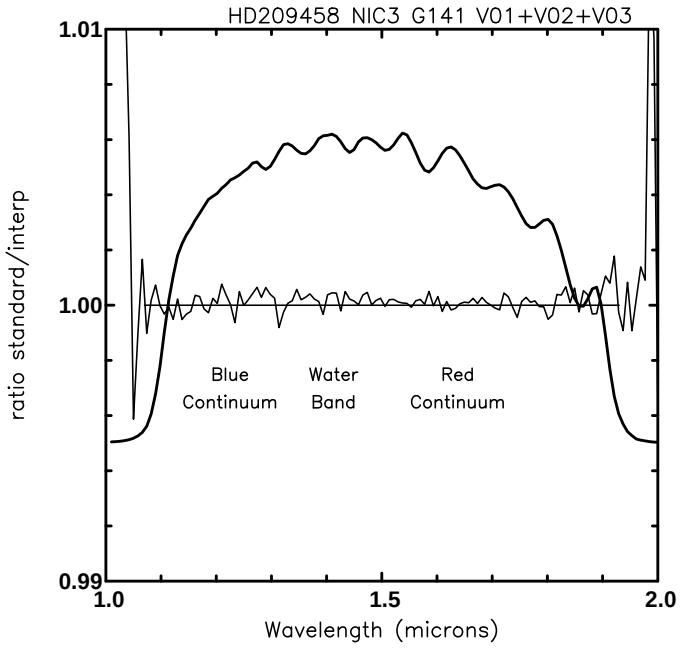

Figure 4: Mean of spectra ratios (multiplied by inverse ratio of exposure times at t_{exp} = 1.82 to 1.99s and 1.99 to 2.11 s plotted as the thin curve. The overplotted heavy curve offset from unity simply illustrates the relative count levels over the grism spectra. There seems to be little if any residual nonlinearity over the spectral domain of interest.

at an 80 s cadence (30% above the Poisson limit for count level of 1.55×10^8 electrons). Rebinned to 80 seconds the NICMOS photometry provides an *rms* of 1.8×10^{-4}, about $\times 2.4$ above the Poisson noise limit (2.16×10^8 electrons detected each 80 s). The noise level on the purely differential index tracking water to continuum bands is closer to being Poisson limited.

6. Summary

With the current loss of STIS on *HST*, the NICMOS grisms provide the best spectroscopic capability for probing atmospheres of extrasolar planets. The presence of broad absorption features, particularly from water vapor, relative to any similar diagnostics in the optical would make NICMOS competitive for current advances, even if STIS had not been lost. The latter of course requires that, like STIS, NICMOS be able to provide data well outside of its design goals in terms of signal-to-noise and stability.

The analyses discussed here have not yet reached as near fundamental precision limits, as was the case for STIS. At a time series signal-to-noise of 5,500 per 80 second sum ($<$ 0.0002 *rms*) for direct photometry summed over the full bandpass, the results are probably far better than could be obtained from the ground, and will be useful for constraining basic system parameters through detailed modeling of the transit light curve. Application of more sophisticated analysis techniques (e.g. difference image analysis) may allow superior results to still be obtained. For the index of water vapor band to continuum counts the current results are already closer to fundamental limits on this very bright star, and should provide interesting constraints on the HD 209458b atmospheric water vapor content.

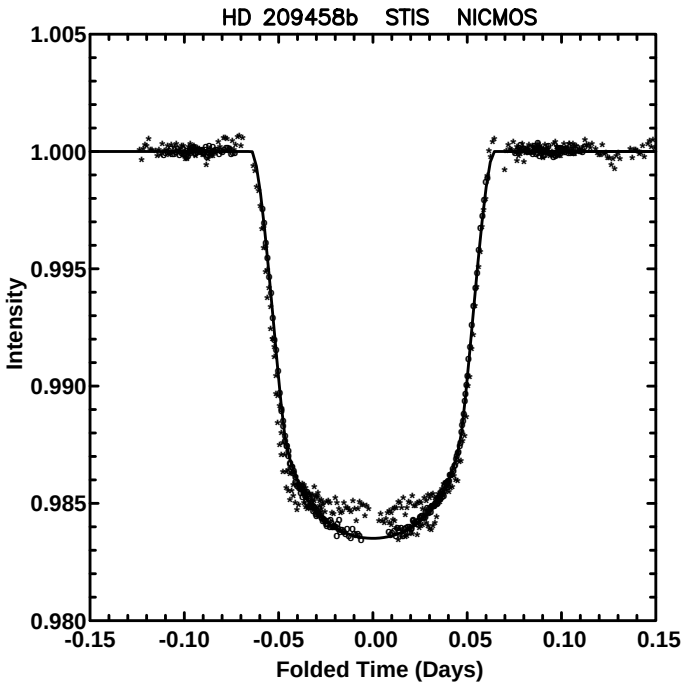

Figure 5: Shows relative intensity through multiple transits from the STIS observations in 2000 'o', and the NICMOS observations in 2003 '*'. The much more rounded bottom to the transit in the STIS (optical) observations is expected from larger limb darkening in the optical than in the IR which should (as generally observed) yield a flatter eclipse bottom.

Acknowledgments. I thank Santiago Arribas, Eddie Bergeron, and Tim Brown for discussion and input on a wide range of topics.

References

Brown, T.M. et al. 2001, ApJ, 552, 699

Brown, T.M. 2001, ApJ, 553, 1006

Charbonneau, D. et al. 2002, ApJ, 568, 377

Gilliland, R.L., and Arribas, S. 2003, *Instrument Science Report* NICMOS 2003-0 (Baltimore:STScI), available through http://www.stsci.edu/hst/nicmos

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

NICMOS Temperature Monitoring

T. Wiklind
[1] *ESA/STScI, 3700 San Martin Dr, Baltimore MD 21218*

T. Wheeler

STScI, 3700 San Martin Dr, Baltimore MD 21218

Abstract. The NICMOS detectors are sensitive to the operating temperature. During the Cycle 7 era, NICMOS was cooled through cryogenic solid nitrogen. This resulted in an operating temperature varying from 61 K to 62.5 K before the nitrogen was exhausted. Using the new NICMOS Cooling System (NCS), the nominal temperature is significantly more stable, but at a higher temperature, 77.15 K. The stability of the dewar temperature is maintained through manual changes to the NCS temperature set-point. The achieved thermal stability is ∼0.1 K. We describe the temperature history, the NICMOS temperature dependencies and remaining temperature fluctuations.

1. Introduction

The Near Infrared Camera, Multi-Object Spectrometer (NICMOS) is the only infrared instrument onboard the Hubble Space Telescope (HST). NICMOS was installed on HST in February 1997 and is described in detail in the NICMOS Instrument Handbook (Schultz et al. 2005). NICMOS thermal history can be divided into two distinct epochs. It was initially cooled to ∼61 K by a cryogenic dewar containing solid nitrogen. The cryogenically cooled dewar suffered a shorter lifetime than expected due to a thermal short, leading to a faster sublimation of the nitrogen. Following nitrogen exhaustion in January 1999, the NICMOS instrument warmed up to a temperature of around 260 K. As this was much too warm for any useful science observations, NICMOS stayed dormant until the third Servicing Mission in March 2002 when a mechanical cooling machine, using a closed-loop reverse-Brayton cycle (Cheng et al. 1998), was installed. The dewar temperature is regulated by the new NICMOS Cooling System (NCS) and is used to maintain the detectors at 77.15 K, about 15 K warmer than with the solid nitrogen cooling.

Operations during the first epoch, with solid nitrogen as the coolant, the detector temperature slowly increased from 61 K to 62.5 K from the start until right before nitrogen exhaustion (see Figure 1). This variation in operating temperature made it necessary to obtain calibration data, i.e. darks, flats and photometric calibration, quite frequently to keep up with the changing thermal environment. During the NCS era, the dewar temperature is higher but also significantly more stable. Temperature variations over the time period March 2002 until the present are less than ∼0.1 K (see Figure 2).

The dewar temperature is presently monitored using the NIC1 mounting cup sensor. This is essentially a diode, with an operating temperature limit of 78 K. With the present nominal operating temperature of 77.15 K, we are fairly close to the upper limit of this

[1] European Space Agency Space Telescope Operation Division

sensor. Should the temperature increase, alternate means of measuring the detector temperatures becomes necessary. One possibility is to use the bias obtained during dark observations. This avenue of monitoring the temperature is presently being explored (Bergeron, in preparation).

2. NICMOS temperature effects

The NICMOS detectors show a number of effects that are sensitive to temperature. It is therefore of utmost importance to maintain the thermal environment as stable as possible over long time scales. Any variation of the detector temperature will require a substantial calibration effort.

The temperature in the dewar depends on several factors, the NCS neon cooling capacity and set-point (the neon temperature in the closed-loop is maintained at ∼72 K and can be adjusted to achieve a desirable detector temperature), and several sources of 'parasitic' heat loading:

- *Annual cycle.* The heat load on HST varies with the Sun distance, with the warmest period in January and the coldest in June. This is seen as a slow change in the average dewar temperature if a constant NCS cooling set-point temperature is maintained.
- *Day/Night Cycle.* Moving from direct sunlight to night changes the heat loading on HST. The time scale is short enough that the NICMOS dewar does not respond quickly. The day/night cycle does, however, provide a thermal driving force which can be detected in the detector temperature.
- *HST Sun/Roll Angle.* The HST sun/roll angle determines which part of the HST aft shroud is in direct sun light. Prolonged pointings can cause more or less heat loading on the NICMOS dewar, with temperature effects of the order 0.1 K on time scales of several days.
- *NICMOS Operation.* Use of the NICMOS detectors, whether for science operations, calibrations or autoreset mode, affect the detector temperature.
- *Other Aft Shroud Instruments.* Instruments other than NICMOS can cause heat loading inside the aft shroud and thereby affect the NICMOS detector temperature on short time scales. This is, however, a relatively small effect.
- *Autoreset Mode.* When not used for science exposures, the NICMOS detectors are continuously read out. The reason for this is to remove trapped charge caused by passages through the South Atlantic Anomaly (see below). This causes a heat loading on the NICMOS detectors, with small (<0.1 K) temperature jitter (see Figure 1).

The temperature response time of the NICMOS detectors is relatively long (∼120 minutes) due to the relatively large thermal mass of the dewar. However, the various thermal loadings listed above induce detectable temperature fluctuations at the NICMOS detectors. In Figure 3 we show three power spectral density (PSD) plots of the NIC1 mounting cup temperature sensor for a period of about two weeks during (a) pre-autoreset mode in Cycle 7, with autoreset mode in Cycle 7, and with the NCS and autoreset mode during Cycle 14. Four distinct periods are distinguishable in all three cases: ∼24 hours, 101 minutes, 49 minutes and 33 minutes. In addition there is a strong peak at ∼3.5 minutes (Figure 4). The latter peak is not related to the temperature of the detector, but is likely to be an artifact caused by the electronics involved in the temperature sensing itself. The 3.5 minute period is present in both pre-NCS and NCS data and is not affected by the autoreset mode (see Figure 4).

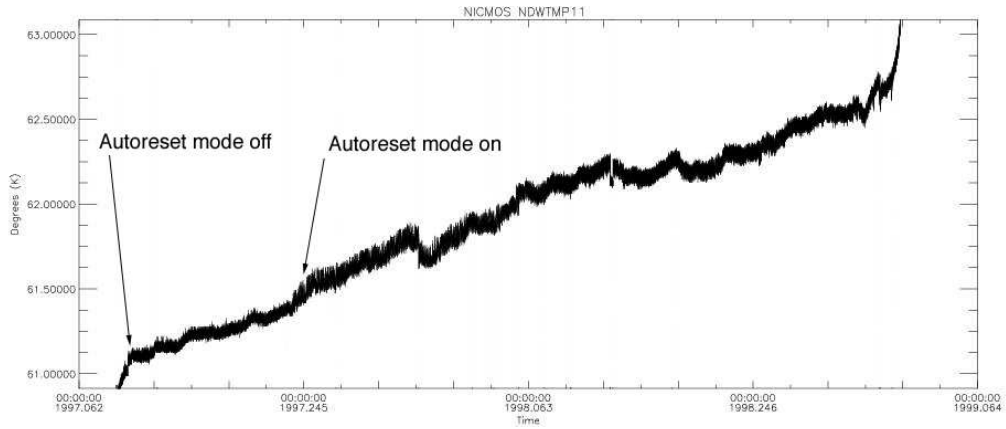

Figure 1: The dewar temperature at the NIC1 detector mounting cup during the period February 1997 until the warm-up in January 1999. An almost linear increase in the temperature, from 61 K to 62.5 K can be seen. The warm-up in January 1999 was caused by the exhaustion of the solid nitrogen supply. At the start of Cycle 7, the autoreset mode (see text) was off. The start of autoreset is seen as an increase in the amplitude of the temperature fluctuations.

The 24 hour period is related to passage through the South Atlantic Anomaly (SAA), which occurs a few times every day and is associated with the deposit of a large amount of charge in the NICMOS detectors. The 101, 49 and 33 minute periods are associated with the orbital period of the HST and, hence, the day/night cycle. Other parasitic thermal loadings are not of periodic nature and hence do not show up in the plot. Although the sampling is too short to show the annual variation, it would be a very weak signal as the setting of the neon temperature set-point is manually regulated to compensate for the slow heating and cooling associated with the seasonal changes.

3. Temperature effects on the NIMCOS detectors

The performance of the NICMOS detectors show several dependencies on temperature. The four most important ones are (1) linear dark current, (2) detector quantum efficiency (DQE), (3) the shading residual and (4) the saturation level.

1. The linear dark current increases linearly with exposure time (hence the name). It is a residual after amplifier glow and shading have been removed. The dark current levels of all three NICMOS cameras are stable, and does not exceed the values expected for the new operating temperature. The increased operating temperature in the NCS era has increased the number of 'hot' pixels, i.e. pixels with higher-than-average linear dark current. This can, however, be accounted for by dithering NICMOS exposures.

2. The detector quantum efficiency (DQE) of the NICMOS detectors depends on the temperature in the sense that the DQE has increased with the new higher operating temperature. Pixels with a lower than average response (i.e. 'hot' pixels) show a higher than average DQE increase with temperature, resulting in a flatter overall response. The average response at 77.15 K increased by about 60% at J, 40% at H, and 20% at K.

3. The NICMOS detectors exhibit a noiseless signal gradient orthogonal to the direction of primary clocking, which is commonly referred to as shading. It is caused by changes of the pixel bias levels as a function of temperature as well as time since the last

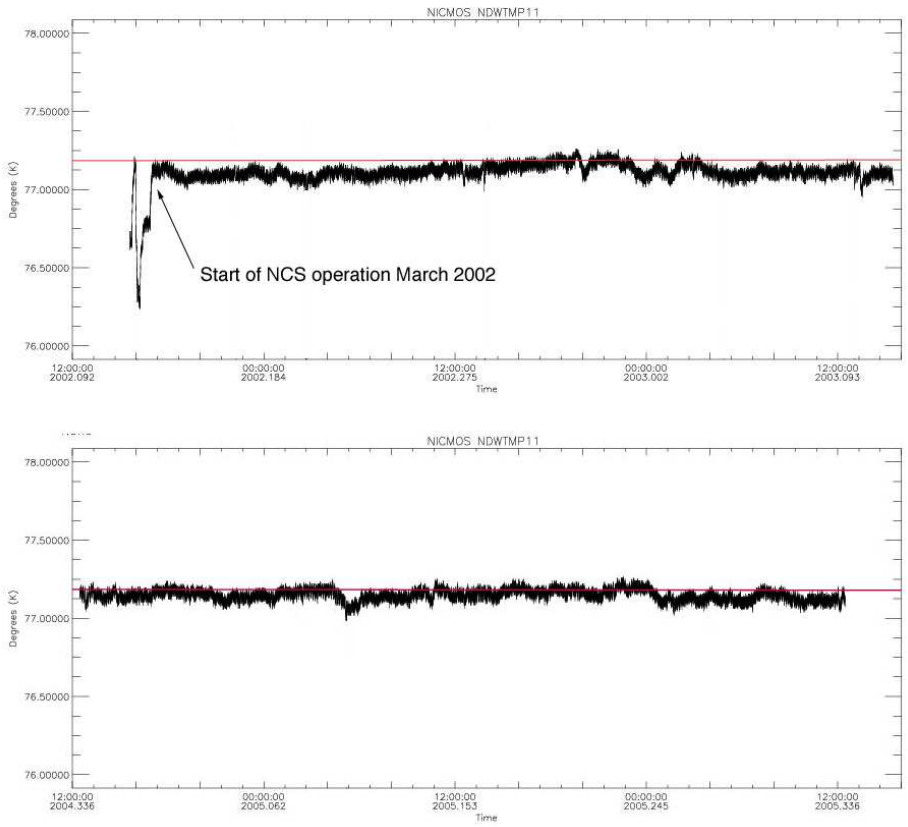

Figure 2: Same as Figure 1 but starting in March 2002 (top figure) when the NICMOS Cooling System was installed. The bottom figure shows the dewar temperature between November 2004 and November 2005. The red horizontal line is the nominal operating temperature of 77.15 K.

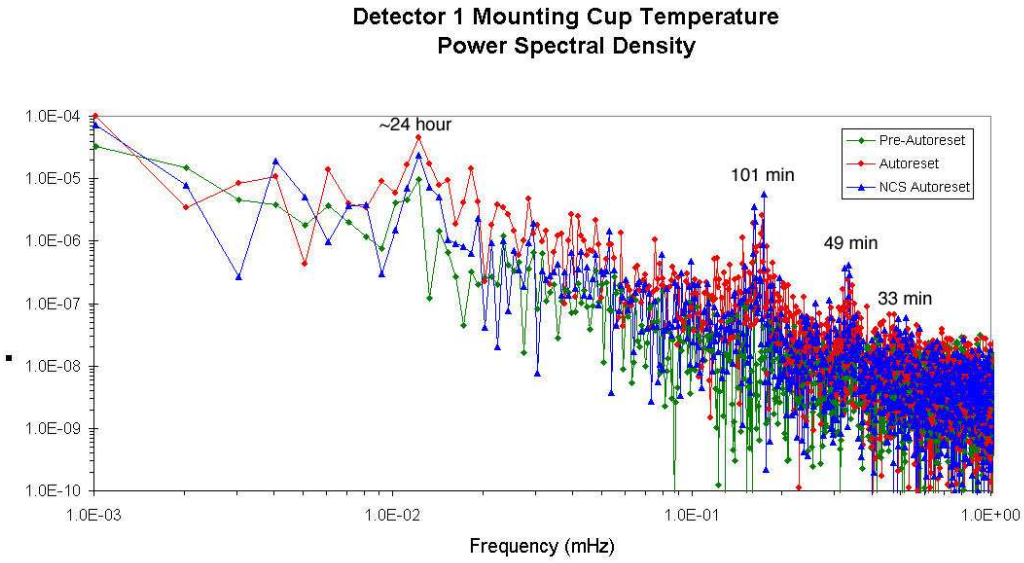

Figure 3: The Power Spectral Density (PSD) of temperature fluctuations at the NIC1 mounting cup for three different thermal epochs: pre-autoreset mode in Cycle 7 (green), autoreset mode in Cycle 7 (red) and NCS with autoreset in Cycle 10-14 (blue). Peaks corresponding to periodicities of ∼24 hours, 101, 49 and 33 minutes are seen in all three epochs. The two first periods correspond to thermal driving from SAA passages and orbital night/day, respectively. The last two periods are harmonics of the orbital periods. Notice that the use of autoreset (autoflush) mode consistently gives a higher power spectral density. This corresponds to the extra thermal load imposed by the continuous reading of the detectors while they are not in use for scientific exposures.

readout (Δ-time). The shading can be completely removed during pipeline processing once it has been calibrated with Δ-time and accurate detector temperature.

4. The saturation level, or dynamic range, of a given detector pixel is defined by the well depth, i.e. the amount of charge loaded into the pixel during the detector reset. The reset voltages of the NICMOS detectors vary with temperature. The pixel saturation level will therefore also vary with temperature. For NIC1 and NIC2, the saturation level is ∼15% lower when operating at 77.15 K than at 61 K, for NIC3, the difference is only ∼7%.

4. Summary

The NICMOS detectors show several effects which are temperature dependent. Although most of these can be handled in the pipeline, successful data reduction requires accurate knowledge of the detector temperature. The easiest way of assuring good data is to keep a constant detector temperature, corresponding to the environment when calibration data was obtained. In the NCS era this is achieved through monitoring of the NIC1 mounting cup temperature sensor and manually changing the NCS neon temperature set-point. Since the re-start of NICMOS in March 2002, we have achieved a temperature stability of ∼0.1 K. Despite this, there are small scale temperature fluctuations on short time scales. These are mainly caused by passages through the SAA, the orbital night/day cycle and intermittent parasitic thermal loads. The autoreset mode, started in Cycle 7 to flush out trapped charges, induces a small scale temperature jitter in the NICMOS detectors.

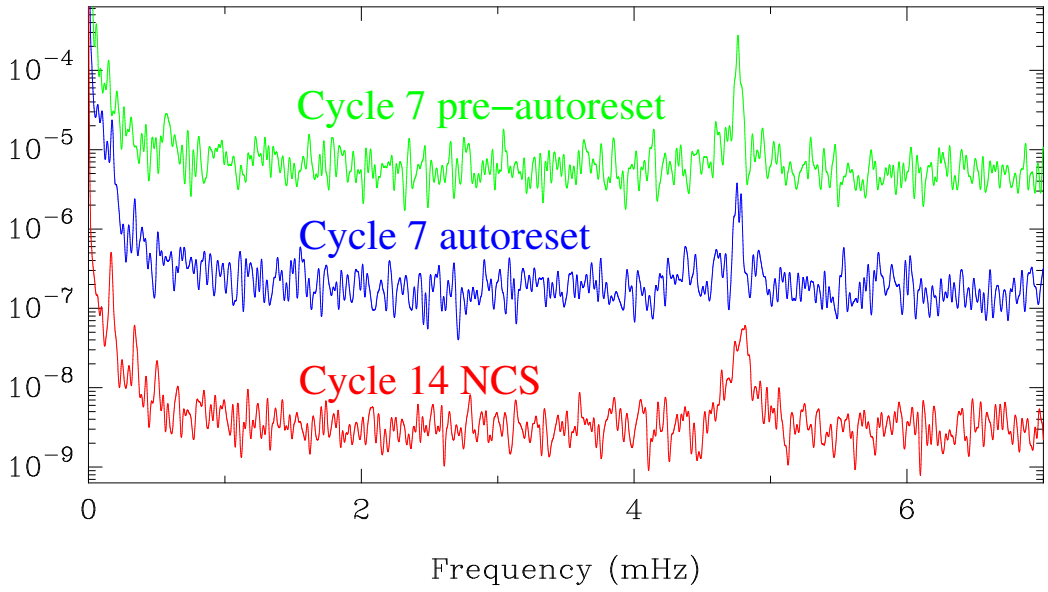

Figure 4: The Power Spectral Density (PSD) of temperature fluctuations at the NIC1 mounting cup for three different thermal epochs: pre-autoreset mode in Cycle 7, post-autoreset mode in Cycle 7 and Cycle 14. The flat part of the frequency response corresponds to white noise and the prominent peak at ∼4.8 mHz (corresponding to a periodicity of ∼3.5 minutes) is not related to thermal fluctuations, but an intrinsic electronic signal. The periods seen in Figure 3 can be seen at the left. The curves have been offset along the y-axis since they would otherwise overlap each other.

References

Cheng, E.S., Smith, R.C., Jedrich, N.M., Gibbon, J.A., Cottingham, D.A., Swift, W.L. & Dame, R.E. 1998, SPIE Proc., 3356, 1149

Schultz, A., Noll, K., Barker, E., Arribas, S., Bergeron L. E., de Jong, R., Malhotra, S., Mobasher, B., Wiklind, T. & Xu, C., 2005, *NICMOS Instrument Handbook*, version 8. (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Removing Post-SAA Persistence in NICMOS Data

E. A. Barker, A. M. Koekemoer and V. Laidler

Space Telescope Science Institute, 3700 San Martin Dr., Baltimore MD 21218, USA

Abstract. A general concern with NICMOS images is the issue of "persistence", or residual charge from bright sources and cosmic rays that remains trapped in pixels and is gradually released during subsequent exposures. This is particularly noticeable after HST transits the South Atlantic Anomaly (SAA), where the cosmic ray rate is so high that their residual flux contributes a significant noise component to many subsequent exposures. We describe the `SAAclean` task which can be run on post-SAA NICMOS images to identify the amount of charge in each pixel that can be attributed to persistent flux from SAA cosmic rays, and applies a correction based on an SAA persistence model to remove this flux from the data. This task has so far been tested on a variety of NICMOS data, and this represents the first delivery of this task in the PyRAF/STSDAS environment to the HST observer community.

1. Introduction

The data taken by HST's Near Infrared Camera and Multi-Object Spectrometer (NICMOS) is particularly susceptible to "image persistence" or latent charge. This residual charge is caused by electrons that become temporarily trapped in the NICMOS detectors. These electrons are gradually released from the detectors in a logarithmic way over time, which produces a latent image of the charge, and can significantly reduce the signal-to-noise (S/N) in subsequent exposures (Bergeron & Dickinson 2003).

Approximately half of all HST orbits pass through the South Atlantic Anomaly (SAA), part of the Van Allen radiation belts. Passing through the SAA produces an extremely large cosmic ray hit-rate, so much so that NICMOS is powered down during the SAA transit. Unfortunately, during this down time, charge accumulates and is trapped in the detectors, so upon powering up the instrument, the residual charge in the detectors is significant. This residual charge escapes slowly, introducing a decaying "persistence" in potentially many science exposures during the same orbit.

Clearly, we would like to recover the lost signal-to-noise in the SAA-impacted science exposures. During Cycle 7, a technique was created (Bergeron & Najita 1998) to create an image of the persistence at the point when the persistence signal was strongest just after exiting the SAA. This "map" of the persistence could then be scaled and subtracted from the subsequent science exposures to remove the persistence and regain much of the lost signal-to-noise.

2. The `SAAclean` Task

In an attempt to automate the process described above, E. Bergeron and M. Dickinson developed the `SAAclean` script, orginally written in IDL. This script is described in detail in NICMOS Instrument Science Report 2003-010 (Bergeron & Dickinson 2003). Here we give a basic overview of the steps involved.

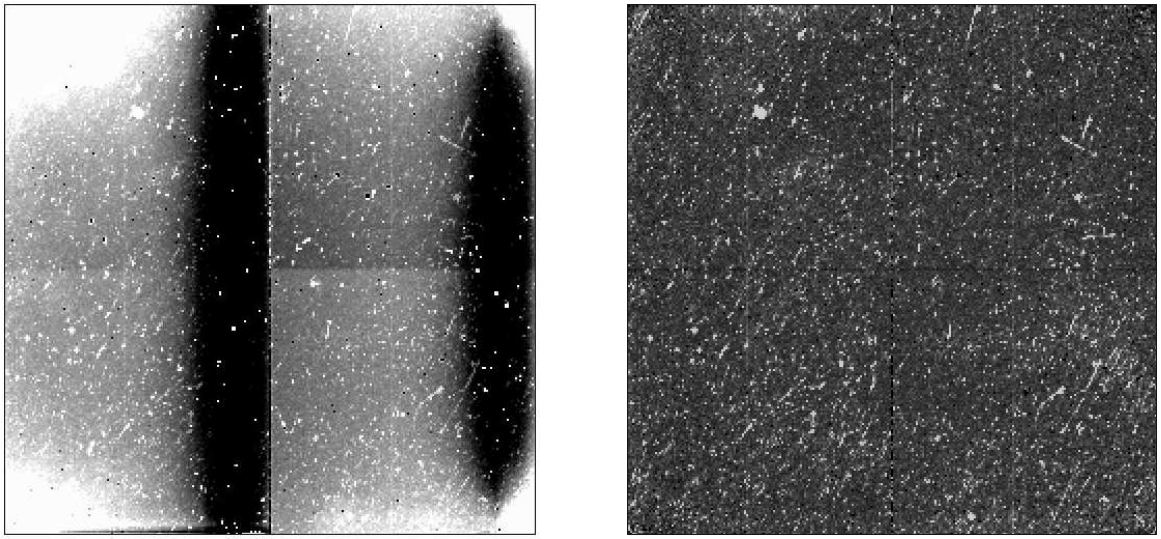

Figure 1: One SAA dark image, scaled to show cosmic rays (left) and the persistence image generated by the `SAAclean` task to clean the SAA-induced cosmic rays (right).

The `SAAclean` script will run through the procedure below if the time elapsed between the science image and the last SAA exit is less than 5600 sec (94 minutes, or 1 orbit) (Bergeron & Dickinson 2003). If the time elapsed is greater than this, the script exits without processing the science image. We presume that science images taken this long after the SAA exit are not affected by the SAA passage. If the time elapsed since the last SAA exit is less than 5600 sec, then the script proceeds as described below. If, after the script has run, the resultant correction to the science image is determined to be less than 1%, the script does not apply the correction and the science image remains unchanged.

It is recommended, but optional, to correct the input science image for the pedestal effect using the `pedsky` or `pedsub` IRAF tasks before processing with the `saaclean` script. We believe performing the pedestal correction will improve the detection of the persistence due to the SAA passage. If a science image has been impacted by the SAA passage, the `SAA_DARK` header keyword will be populated with the name of the association table that points to the two post-SAA dark images correlated with the science image. Using the two post-SAA dark images, a decay scale factor is determined, and then the two images are averaged together with some cosmic ray rejection. The result is then pedestal corrected, which provides an image of the persistence (see Figure 1).

This image of the persistence is then iteratively scaled by the scale factor between the two post-SAA dark images, fitting to the scale factor that provides the minimum total noise in the final output image, and then subtracted from the input science image. If desired, an additional and optional pedestal removal can be done at this point. We expect the results of this pedestal correction to be better than the first pedestal correction, since the persistence has now been removed from the science image.

Bergeron's original IDL script has been converted to a PyRAF task in Python code by V. Laidler at STScI. As with any such conversion between languages, some algorthmic changes are occasionally necessary. These algorthimic changes are kept to an absolute minimum and used only when the original code does not have a precise equavalent in the new language.

The `SAAclean` task produces several intermediate and final statistics as output data, indicating the values it calculates in determining a possible correction value to be applied to the science data (see the example output in Figure 2). This output includes the `SAAclean` input filename, persistence scale factor, threshold for determining the high S/N pixel regime

```
saaclean version 0.5 (15 Sep 2005)
Input file: n90j55z2q_ped.fits
Using scale factor of  0.54  to construct persistence image
Bad pixel image filename obtained from  /data/cdbs5/nref//m991700on_flt.fits
median used in flatfielding:  0.0304204884365
Threshold for hi/lo:  0.0698112137616
Npixels hi/lo:  8000 57536

Results summary for high domain:
chi2 for parabola fit =  0.00074004532324
   min-noise (best) scale factor is:  0.250980787999
   effective noise at this factor (electrons at gain 5.400000): 41.598116
   noise reduction (percent)       :  15.945775261

Results summary for low domain:
chi2 for parabola fit =  2.1850359789e-05
   min-noise (best) scale factor is:  0.216667701363
   effective noise at this factor (electrons at gain 5.400000): 32.178225
   noise reduction (percent)       :  2.4951591989

 Applying noise reduction in both domains
```

Figure 2: An example of `SAAclean` output.

versus the low S/N pixel regime, the number of pixels in each regime, and best fit factors and noise values for each regime.

As noted in the original `SAAclean` ISR (Bergeron & Dickinson 2003), any correction determined by the `SAAclean` task is applied to the input science exposure (see Figure 3). We assume here that the given recommendation of pedestal correcting the science exposure before running `SAAclean` has been done. If so, this implies that the `SAAclean` correction is applied to the pedestal corrected science image.

3. Testing of the `SAAclean` Task

We have tested the `SAAclean` task on the vast majority of existing NICMOS data, utilizing all three cameras (NIC1, NIC2, NIC3). In order for `SAAclean` to run, several header keywords must be populated in the input science exposure header. The `FLATFILE` keyword must be populated with the name of the corresponding flat field image filename. The `SAA_DARK` keyword must be populated with the name of the association table containing the names of SAA dark images associated with the science exposure. These keywords are typically populated automatically by the OTFR processing when the data is requested from the archive.

We are currently making slight modifications to the `SAAclean` task in an aim to improve functionality over a broad range of data types and to make it slightly more user friendly. We have found that enabling the use of lists of filenames (@-files) can make the task more more efficient to use. We have also modified the original script behavior of using an input parameter to determine the gain of the science image. We are now extracting the gain given in the science image header. In addition, the task uses the bad pixel file specified in the image header to flag pixels that should not be used to determine the SAA correction.

Finally, the Pyraf task differs from the IDL version in that, while it still calculates the SAA image using a pedestal-corrected image, the resulting correction is applied directly to the calibrated NICMOS exposure (which has not necessarily been run through pedsky). This provides users with more control over the processing steps, since it is often desirable to run additional pedestal correction after having applied the SAA correction.

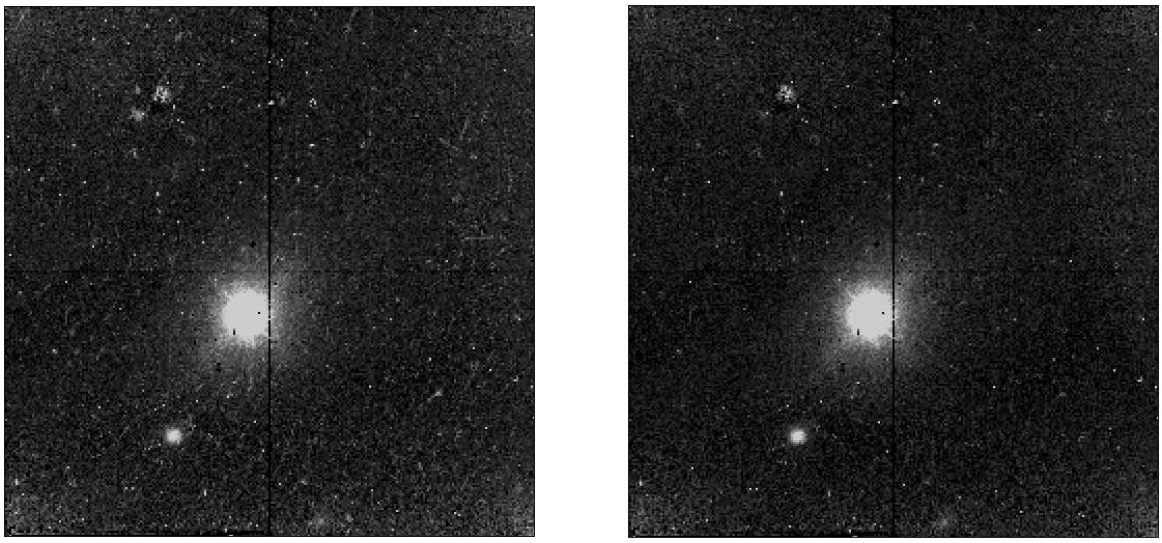

Figure 3: An SAA-impacted, input pedestal-corrected image (left) and the output `SAAclean` image (right), which has fewer low-level residual cosmic rays.

4. Summary

NICMOS images are strongly impacted by SAA passages through latent charge remaining in the detectors, which causes "persistence" in subsequent exposures. We have tested a new PyRAF task, based on the `SAAclean` script developed by Bergeron and Dickinson, which can remove much of the SAA-induced persistence and regain otherwise lost signal-to-noise.

After testing the new `SAAclean` task on essentially all existing NICMOS archived data, we see that in most cases, `SAAclean` generally improves the signal-to-noise of objects of interest, unless there are problems with the data unrelated to the SAA passage.

5. Future Work

We are currently in the process of modifying the task to be incorporated into the standard calibration pipeline for NICMOS data. The pipeline task would be performed after `calnica` and before `calnicb`. We do not expect to change the underlying `SAAclean` algorithm, but various keyword functions and properties are being handled such that they are pipeline-standard compliant. We expect the pipeline version of the `SAAclean` task to be available to the community by late spring or early summer of 2006.

References

Bergeron, L. E. & Dickinson, M. E. 2003, *Instrument Science Report* NICMOS 2003-010 (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Bergeron, L. E. & Najita, J. R. 1998, BAAS, 30, 1299

Part 3. STIS, FOS and GHRS

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The STIS Closeout Plan

Paul Goudfrooij and the STIS Team: Alessandra Aloisi[1], Rosa Diaz-Miller, Linda Dressel, Jessica Kim Quijano, Jesús Maíz-Apellániz[1], and Charles Proffitt[2]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. The huge archive of STIS observations constitutes a unique —and arguably the most comprehensive— source of spectroscopy in the UV and spatially-resolved spectroscopy in the optical available to the entire astronomy community. The main goal of the STIS close-out plan is to render the full archive of STIS data in a calibration status suitable for performing high-level and high-accuracy science before STIS support at STScI is phased out. In addition, we intend to implement software improvements as well as enhancements related to the retrieval of STIS data from the *HST* archive. Detailed information on the various close-out activities is presented.

1. Introduction

The Space Telescope Imaging Spectrograph (STIS) stopped science operations on August 3, 2004, due to a failed power supply in its side-2 electronics (its side-1 failed already in May 2001; more details on the STIS failure can be found in Goudfrooij 2005). STIS was an extremely versatile spectrograph and imager offering a diversity of instrument modes unmatched by any other *HST* instrument. Since the start of Cycle 7 (the first *HST* cycle in which STIS was available), STIS was one of the most productive *HST* instruments. As of November 1, 2005, over 900 refereed publications using STIS have appeared in the literature[3], with a total number of citations exceeding 16,000. Analysis of STIS data has had a major impact on virtually all science categories, in particular the morphology and physics of stellar outflows, the composition and dynamics of the interstellar medium, the demography of super massive black holes in the centers of galaxies, and the baryon content in the local universe.

It is clear that the huge archive of STIS observations constitutes a unique and comprehensive source of UV spectroscopy and spatially-resolved spectroscopy in the optical region, available to the astronomy community. As such, it deserves an appropriately high level of attention from STScI in terms of providing adequate calibrations before STIS team resources are taken away. The STIS Closeout Plan was designed with this in mind.

2. Main Goals of STIS Closeout Plan

The extraordinary versatility of STIS presented a significant challenge to the STIS team in terms of being able to keep up with the characterization and calibration of all sup-

[1]Affiliated with the Space Telescope Division, European Space Agency

[2]Operations Astronomer, Computer Science Corporation

[3]Using information from the NASA Astrophysics Data System (ADS)

ported observing modes[4], many aspects of which are/were complex and/or time-dependent. The main goal of the STIS Closeout Plan is to render the archive of STIS data taken in 'supported' modes (plus a subset of the 'available-but-not-supported' observing modes; see Table 1) in a calibration status suitable for performing high-level and high-accuracy science before STIS support at STScI will have to be phased out for the benefit of new instruments to be installed in *HST* Servicing Mission 4. We currently estimate that this phase-out for STIS support will occur by the end of 2006.

The goal for the final overall level of calibration accuracy to be reached for all (previously) supported STIS observing modes is listed in Chapter 16 of the STIS Instrument Handbook (Kim Quijano et al. 2004). In addition, we intend to implement software improvements as well as enhancements related to the retrieval of STIS data from the *HST* archive. The activities in the STIS Closeout Plan can be divided into one of the following three main classes:

1. Already planned tasks of moderate-to-high priority which did not yet come to completion. These tasks are related to STIS calibration of regularly used observing modes and the subsequent implementation into software (within the pipeline or as off-line IRAF tools) or pipeline reference files;

2. New tasks related to, e.g., the enhancement of archive capabilities regarding STIS data, or to a proper documentation of all activities regarding the use of STIS data and of the monitoring activities used by the STIS branch over the years;

3. Tasks related to the implementation by STScI of the software and other deliverables created by the Space Telescope-European Coordinating Facility (ST-ECF) in Germany as part of its STIS-CE (Calibration Enhancement) project.

In the following sections we will provide the reader with specific, but brief, information about each ongoing or planned significant activity in the STIS Closeout Plan as of Nov 15, 2005. The activities are organized by main category: Calibration, Pipeline, Archive Enhancement, and User Information. The overall priority of each activity is indicated after each title; the priorities were assigned based mainly on two factors regarding the observing mode involved, namely *(i)* statistics on its usage, and *(ii)* the uniqueness of its capability relative to those offered by other instruments. For the calibration activities, we also list the ID of the calibration program(s) involved if appropriate.

3. Description of Activities in the Closeout Plan

3.1. Calibration Activities

E-Mode Sensitivities (Priority: Crucial): This activity will deliver the final sensitivity calibrations for the echelle modes of STIS. This includes final reference file updates for wavelength-dependent aperture corrections, wavelength-dependent extraction size corrections, and wavelength-dependent time dependences, using all calibration data collected through the summer of 2004. The SYNPHOT component files for the echelle modes also need to be updated (the current versions are not consistent with the photometric sensitivity reference files in the pipeline; e.g., the former do not include the echelle blaze functions, only the 'envelope' of the sensitivity curves). Results of this activity will be documented in an Instrument Science Report (hereafter ISR, which are available through the STIS instrument web site at http://www.stsci.edu/hst/stis). (See Aloisi (2006) in this volume.)

[4]The supported set of observing modes included: 3 *detectors* (FUV-MAMA, NUV-MAMA, CCD), 4 operational *modes* (ACQ, ACQ/PEAK, ACCUM, TIMETAG), 44 *apertures* (filters, slits), and 133 grating/tilt combinations.

Table 1: Status of and calibration plan for spectroscopic STIS observations taken in "available-but-not-supported" mode[a]

Observing Mode	# exp.	Calibration Plan or limitation of data
Non-recommended echelle grating + small aperture combination	21	To be calibrated for final STIS archive
Echelle and first-order spectra taken with non-supported aperture	86	Limitations for each grating/aperture combination will be indicated in final close-out report
First-order spectra taken with 2×2 or 6×6 aperture	2587	Planning to calibrate for final archive (treated as slitless for 6×6 aperture).
Spectra taken with unsupported fiducial bars	67	Limitations for each grating/aperture combination will be indicated in final close-out report
Echelle spectra of extended objects with long slits	39	Calibration needs to be different for every individual science application. Will stay unsupported.
Spectra in G750L @ 8975Å or G750M @ 10363Å	3	2^{nd}-order overlap. Will stay unsupported; limitations will be indicated in final report.

[a] There are 3666 spectroscopic science observations in the STIS archive taken in "available" mode (header keyword CFSTATUS = 'AVAILABLE'), which is ∼10% of the total.

HST Program Numbers: 8915, 8919, 9628, 10033

E-Mode Sensitivity Uncertainty due to Blaze Shifts (Priority: High): This activity is related to the previous one but is listed separately because it includes the complicated effort to derive a correction for the so-called Monthly Offsets and their associated blaze shifts, which do not repeat well for a given month of the year (i.e., the time dependence of the amplitude of the blaze shifts is currently not adequately corrected for). This effort shall be performed for both primary and secondary central wavelength settings, as they have had similar usage by GOs. This activity will involve writing at least one ISR, and may involve the need for code changes within CALSTIS. (See Aloisi (2006) in this volume.)

HST Program Numbers: 8915, 8919, 9619, 9628, 10033

MAMA Full-Field Sensitivity Monitor (Priority: Medium): These activities comprise flux measurements of various isolated stars covering the field of view of the STIS MAMA detectors, from images that are taken at regular time intervals so that the same star is imaged on different parts of the detector. These data are used to monitor the full-field sensitivity of the detectors as well as the astrometric and PSF stabilities. The data from Cycles 11 and 12 will have to be analyzed. The results will be compared to the previously determined time dependence of the sensitivity for the different imaging modes, and to evaluate the need for updates to the low spatial frequency flat fields for imaging modes. This activity will involve writing one ISR.

HST Program Numbers: 9622, 10028, 9623, 10032

MAMA Dark Monitor (Priority: Medium/High): This activity involves the measurement of the dark current levels of the two MAMA detectors, which have been

taken twice a week for each detector. The NUV-MAMA dark data taken over the years will be assembled together to determine final functional forms of the global dark current levels as a function of time, tube temperature, and charge amplifier temperature. These scaling functions will be implemented into the STIS calibration pipeline (CALSTIS), involving new reference files as well as a code change. In the case of the FUV-MAMA detector, additional work is needed to determine a better parametrization of the dark level in the "glow" region of the detector. We already have a good idea on this from earlier preparation work, namely to parametrize the FUV-MAMA dark level as a function of the time elapsed since the high-voltage was turned on. For a "heritage" instrument, this can be implemented by delivering a reference table containing the turn-on times, or as a post-observation tool using another proxy that is dependent on the time elapsed since the high-voltage was turned on, e.g., the intensity level of hot pixels. These findings will be written up as two ISRs.

HST Program Numbers: 9615, 10034

CCD Bias and Dark Monitor (Priority: High): This activity involves a thorough investigation of the evolution of the CCD biases and darks as a function of time, including the number of hot columns in superbiases, the level of the spurious change in superbiases (Goudfrooij & Walsh 1997), and the number of hot pixels in darks at different intensity levels and as a function of row number (indicating the evolution of CTE loss). The main purposes of this activity are to provide useful information on the behavior of the performance of CCDs (of different architecture) in orbit, and to compare the results with the performance of other CCDs on HST. This activity will likely involve writing two ISRs.

CCD External Sparse-Field CTE monitor (Priority: High): This activity involves the analysis of the last two epochs of the *External* Sparse-Field CTE monitor calibration program. These data bear directly on the influence of CTE effects to the accuracy of point source photometry and spectroscopy. The analysis has been described before by Goudfrooij & Kimble (2002). The results will also be compared to those derived from the physical model of CTE loss of the STIS CCD by P. Bristow of the ST-ECF (see below). This activity will involve writing two ISRs, and it will also be used in the final version of a manuscript on the CTE performance of the STIS CCD, to be submitted to PASP.

CCD CTE effect on Extended Sources (Priority: Med/High): This activity involves the analysis of data observed during three calibration programs to determine the effects of CTE loss to surface photometry and spectroscopy of extended sources, as well as the time dependence of those effects: *(i)* the effects of CTE loss to CCD imaging of galaxies (e.g., luminosity, ellipticity, and position angle measurements at a given surface brightness), *(ii)* The effect of CTE loss to continuum fluxes and emission-line intensities in spectroscopy mode; *(iii)* The effect of CTE loss to *absorption*-line intensities, profiles, and equivalent widths in spectra. All results mentioned above will be compared to the correction applied by the physical model of CTE loss of the STIS CCD by P. Bristow of the ST-ECF (see below). In addition, the results from *(ii)* and *(iii)* above will also be compared to the CTE correction provided by the STIS pipeline (derived from point source spectrophotometry). This activity will involve writing three ISRs.

HST Program Numbers: 8839, 8927, 10038

Imaging Zeropoints and Color Terms (Priority: High): This activity involves the analysis of imaging data of sources with known photometry and covering a wide range of *(known)* SEDs (spectral energy distributions). The analysis will involve

deriving the influence of the intrinsic colors to the derived zeropoints (i.e., the so-called 'color terms' in the photometric calibration), and a comparison with predictions using synthetic spectra (using SYNPHOT). This activity will involve writing a PASP paper. (See Proffitt (2006) in this volume.)

MAMA First-Order Dispersion Solutions (Priority: High): This activity involves the analysis of deep first-order MAMA wavecals taken in the context of a calibration program. Internal wavecals have been obtained at all primary and secondary central wavelengths. Exposure times were chosen to yield enough strong emission lines to constrain adequate wavelength solutions. Data were taken at the zero MSM offset position which is in the middle of the range covered by monthly offsets, and hence provides the best average dispersion solution. Dispersion solutions will now be derived using the recently published Pt/Cr-Ne line list (Sansonetti et al. 2004). (Note that the current dispersion solutions were derived using a Pt-Ne line list, even though the line lamps used on STIS were Pt-Cr-Ne hollow cathode lamps). A new dispersion (_dsp) reference file will be created and delivered. An ISR describing dispersion changes as well as the accuracy of the previously used dispersion solutions will be written and published.

HST Program Numbers: 9618

Grating Scatter for the G230LB Grating (Priority: Medium): This activity involves the analysis of measurements of red targets taken with *both* the G230L grating (a NUV-MAMA mode) *and* the G230LB grating (a CCD mode). The goal is to determine a correction for the influence of grating scatter (primarily from the far wings of the LSF, similar to the effects seen before for the UV gratings used in the Faint Object Spectrograph (Rosa 1994)) in case of CCD/G230LB spectra of red targets. This activity will involve communication with M. Gregg (see Gregg et al. (2006) in this volume) who is working on this issue as well in the context of his STIS SNAP program.

HST Program Numbers: 7723

Spectroscopic PSF Across Slit (Priority: High): This activity involves the characterization of the PSF across the slit for two commonly used long slits (52x0.1 and 52x0.2). Multiple G750L spectra are taken of a K giant star, stepping the slit across the star between exposures to sample the PSF along the dispersion direction. For chosen wavelength intervals, the relative fluxes in each slit position and spectral row are compared to the values predicted using TinyTim models. It is important to verify these models since they are used as input in the dynamical modeling of spectral images, e.g., in the dynamical modeling of galactic nuclei. This activity will involve writing one ISR. (See Dressel (2006) in this volume.)

HST Program Numbers: 9610

Faint Standards Extension (Priority: High): This activity involves analysis of CCD spectra of (faint) white dwarfs which were previously identified and verified as *bona fide* spectrophotometric standard stars. The purpose of this analysis is to provide a thorough verification of the previously established CTE corrections for CCD spectroscopy, namely by stepping the target along the slit (5 positions) with two (short) exposure times. This will verify the results using the two-amplifier readout method, and provide high-S/N data at low intensity levels and low background level. This activity involves the writing of one ISR and may also involve an update of the CTE parameters in the CCD Table Reference File, if necessary. First results from this activity are shown in this volume (Goudfrooij & Bohlin 2006).

HST Program Numbers: 10037, 10039

Trace Stability for Often-Used Modes (Priority: High): Spectral "trace" reference files prescribe the projection of spectra onto the detector at a given position along the spectrograph slit. They are used to produce rectified spectral images and extracted spectra. Accurate traces are needed when individual rows in spectral images are to be analyzed – e.g., in kinematic studies of galaxies, designed to measure the masses of supermassive black holes in galactic nuclei. Accurate traces are also needed for photometric accuracy when small extraction boxes are required to separate nearby point sources or to better isolate a point source from more extended emission. The spectral traces now in use were created from inflight data taken early in the orbital lifetime of STIS, and they can be seen to be in error in rectified spectral images taken in the last few cycles. This project will use calibration as well as GO data to derive spectral traces near the center of the detector (and, for CCD observations, at the E1 aperture positions), for the most commonly used gratings and central wavelengths at several epochs. We will examine changes in rotation and shape of traces, and produce new reference files where needed. The highest-priority grating settings, selected by total science observing time, are all L gratings as well as G750M at central wavelengths 6581, 6768, and 8561. Apart from the creation of calibration reference files where needed, this project will involve the writing of an ISR. (See Dressel et al. (2006) in this volume.)

3.2. Pipeline Activities

Rectification of Non-Dithered Spectra of Spatially Resolved Targets (Priority: High): This activity is to improve the quality of 1 – 3 pixel high extractions of STIS CCD spectra taken with gratings for which the traces are tilted significantly with respect to the pixel array (especially G430M and G750M). The current pipeline extractions use bilinear interpolation which produces strong undulations in spectral extractions of 1 – 3 pixels high. This activity is expected to involve interactions with the Eta Car Treasury Program[5] who have already implemented an interpolation routine which seems to work much better in that respect than that in our STIS pipeline. This activity will also involve writing an ISR. (See Davidson (2006), Barrett and Dressel (2006), and Dressel et al. (2006) in this volume.)

Spectral Dithering within MultiDrizzle (Priority: High): This activity will allow the handling of spectroscopic STIS data within the MultiDrizzle tool. The intent is to allow combination of sets of spectral images which involve inter-exposure offsets (dithers) in the dispersion direction (across the slit), the spatial direction (along the slit), or both. The correction for dithers in the dispersion direction will involve calibrations that are dependent on grating and aperture. We anticipate that this activity will involve a significant amount of testing by STIS Team members. It will also involve writing an ISR and —likely— the need for delivery of reference files that allow for a correction of the throughput loss associated with stepping a source across the slit.

STScI/ST-ECF Collaboration Items (Priority: High): This activity involves STIS Team work related to the implementation of deliverables of the STIS Calibration Enhancement (STIS-CE) project within the STScI environment, be it decision making and/or providing assistance regarding software modules within the CALSTIS pipeline, the STIS archive, or STIS-specific software to be released as IRAF/STSDAS tasks. The two main remaining projects related to ST-ECF deliveries will constitute *(i)* the testing and possible implementation of a new, physical model-based wavelength calibration module for the high-resolution echelle gratings, and *(ii)* the testing, evaluation, and possible implementation of the physical model-based correction for CTE

[5]see http://etacar.umn.edu

loss of the STIS CCD by P. Bristow. (See Bristow et al. (2006) and Kerber et al. (2006) in this volume.)

Final Calibration of STIS Data (Priority: High) This activity involves a comprehensive run of all archival STIS data through OTFR once all final calibrations, pipeline coding, and reference files have been tested and delivered to the various databases. The final set of calibrated STIS data will then be stored in the *HST* archive, and OTFR can be switched off for STIS use which would make more processing power available for OTFR requests of active *HST* instruments. It would also render calibrated STIS data to be readily available for viewing and plotting, which is relevant in the Virtual Observatory era.

3.3. Archive Enhancement Activities

Imaging PSF Library (Priority: Medium): This activity involves the work needed to allow the user to retrieve well-exposed imaging PSFs in a given observing mode (detector / filter combination) and location on the detector. This consists mainly of the identification and assembly of a data base of appropriate datasets, the creation and implementation of a Graphic User Interface (GUI) that performs the archive retrieval of the datasets involved, and the writing of an ISR (which can in principle be combined with the next activity). It can be foreseen that the GUI mentioned above would be developed for *all HST* instruments simultaneously.

Spectroscopic PSF Library (Priority: Medium/High): This activity involves work needed to allow the user to retrieve well-exposed spectroscopic PSFs in a given observing mode (detector / grating / slit combination) and location on the detector. This consists mainly of the identification and assembly of a data base of appropriate datasets, the creation and implementation of a Graphic User Interface that performs the archive retrieval of the datasets involved, and the writing of an ISR (which can in principle be combined with the previous activity). It can be foreseen that the GUI mentioned above would be developed for *all HST* instruments simultaneously.

GO Wavecal Association (Priority: High): This activity involves the development of a system (to be implemented within the *HST* archive) that automatically associates STIS GO wavecals (i.e., wavecals inserted by STIS GO's in their Phase-II proposals, which typically happens when they were allowed to forego the default automatic insertion of wavecals) with the appropriate science spectra, so that retrieval of the latter data will automatically attach the GO wavecals to them. (Currently, only science data are returned *(no wavelength calibration is performed by the pipeline)*, and the GO will need to issue a second archive query to find the GO wavecals based on the proposal/visit combination of the science data. Note also that if a calibrated spectrum is made without a proper wavecal, then serious errors in the wavelength and flux scales can result.) It can be foreseen that this activity would be performed in concert with association-related activities for other *HST* instruments.

Fringe Flat Association (Priority: High): This activity involves the development and implementation of a system within OPUS and the archive that automatically associates contemporaneous STIS CCD Fringe Flats (which have been inserted by STIS GO's in their Phase-II proposals when G750L or G750M spectra were taken) with the appropriate science spectra, so that retrieval of the latter data will automatically attach the appropriate fringe flats to them. (Currently, only the science data are returned and the GO will need to issue a second archive query to find the contemporaneous fringe flats based on the proposal/visit combination of the science data.) It can be foreseen that this activity would be performed in concert with association-related activities for other *HST* instruments.

Spectroscopic Preview Enhancement (Priority: High): This activity involves a review and improvement of the "preview" facility available within the *HST* archive pages for STIS spectra (the "Plot marked spectra" button), whose main purpose is to simplify the preparation of archival studies or proposals. While this facility certainly provides useful output already, there are several aspects that can be improved significantly, e.g., in the areas of the sky subtraction method currently in use and the automatic assignment of display parameters that are often imperfect. Moreover, the facility should also be able to present a 2-D image of the spectral data, so users can see the spatial extent of the target.

3.4. User Information Activities

Data Handbook Update (Priority: High): This activity involves a final review and update of the STIS chapter of the *HST* Data Handbook. The last update was made before a significant number of updates to the CALSTIS pipeline and stand-alone tasks in the STIS package within IRAF/STSDAS were implemented. Hence it is important to create a final, all-encompassing version of the Data Handbook for STIS. It should include a data analysis "cookbook" to guide users in the routine analysis of STIS spectral data.

Summary Document: The STIS Experience (Priority: High): This activity involves the writing of a document that summarizes our experience with the operation and calibration of the instrument. It will include sections on the MAMAs, the CCD, the Optics, etc. There will be many references to other available reports (e.g., ISRs, Instrument Handbook), but all kept in one document to provide an easy reference for comparisons with the operation and calibration of similar detectors in other (present and future) *HST* instruments.

Acknowledgments. The remarkable science legacy of STIS is the tangible result of years and years of effort by many individuals in the STIS Instrument Definition Team, Ball Aerospace, Goddard Space Flight Center, STScI, and elsewhere who defined, built, improved, tested, characterized, calibrated, and explained the instrument. In closing, we wish to thank the following individuals[6] who contributed so much of their time, effort and talent over the last 15 years for the benefit of STIS users. They have made working on STIS an amazing experience.

STIS Instrument Definition Team (IDT): B. Woodgate, T. Gull, A. Boggess, C. Bowers, A. Danks, R. Green, S. Heap, J. Hutchings, E. Jenkins, C. Joseph, M. Kaiser, R. Kimble, S. Kraemer, J. Linsky, S. Maran, H. Moos, F. Roesler, G. Timothy, D. Weistrop

STIS IDT Contractors: T. Beck, N. Collins, J. Feggans, W. Fowler, C. Grady, R. J. Hill, R. S. Hill, W. Landsman, D. Lindler, E. Malumuth, P. Plait, R. Robinson, J. Sandoval, C. Standley

STScI STIS Team, current: A. Aloisi, R. Diaz-Miller, L. Dressel, P. Goudfrooij, J. Kim Quijano, J. Maíz-Appellániz, C. Proffitt

[6]We know we are taking a risk of missing someone who contributed to STIS development. We apologize to those missed in this list.

STScI STIS Team, ex officio: S. Baum, R, Bohlin, T. Brown, J. Christensen, M. Clampin, I. Dashevsky, J. Davies, R. Downes, B. Espey, H. Ferguson, A. Gonnella, S. Friedman, G. Hartig, J. Hayes, S. Hulbert, R. Katsanis, S. Keener, E. Kinney, G. Kriss, C. Leitherer, M. McGrath, G. Mitchell, B. Mobasher, K. Sahu, A. Schultz, R. Shaw, E. Smith, D. Stys, J. Valenti, N. Walborn, J. Wilson

ST-ECF STIS Support: A. Alexov, P. Bristow, M. Fiorentino, F. Kerber, M. Rosa, J. Walsh

STScI STIS Engineering and Software Support: V. Balzano, P. Barrett, I. Busko, P. Hodge, C. Long, R. Pitts, T. Wheeler, D. Zak

Ball Aerospace: W. Meyer, D. Hood, V. Argabright, S. Becker, M. Bottema, R. Breyer, R. Bybee, P. Christon, A. Delamere, D. Dorn, S. Downey, P. Driggers, D. Ebbets, J. Gallegos, H. Garner, J. Hetlinger, R. Lettieri, C. Ludtke, D. Michika, R. Nyquist, D. Rose, R. Stocker, J. Sullivan, C. van Houten, R. Woodruff

References

Aloisis, A. 2006, this volume, 190

Barrett, P. & Dressel, L. 2006, this volume, 260

Bristow, P. et al. 2006, this volume, 299

Davidson, K. 2006, this volume, 247

Dressel, L. 2006, this volume, 277

Dressel, L. et al. 2006, this volume, 267

Goudfrooij, P., 2005, *STScI Newsletter*, 22, 6

Goudfrooij, P., & Walsh, J., 1997, *Instrument Science Report* STIS 1997-09 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

Goudfrooij, P., & Kimble, R. A., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 105

Goudfrooij, P., & Bohlin, R. C., 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 289

Kim Quijano, J., et al., 2004, *STIS Instrument Handbook* (Baltimore: STScI)

Gregg, M.D. et al. 2006, this volume, 209

Kerber, F. et al. 2006, this volume, 309

Proffitt, C. R. 2006, this volume, 234

Rosa, M. R., 1994, in Proc. 1993 *HST Calibration Workshop*, ed. J. C. Blades & S. J. Osmer (Baltimore: STScI), p. 190

Sansonetti, C. J., Kerber, F., Reader, J., & Rosa, M. R., 2004, ApJS, 153, 555

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Towards a Comprehensive Sensitivity Calibration of the STIS Echelle Modes

Alessandra Aloisi[1]

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. I will present the state of the art of the final sensitivity calibration of the STIS spectroscopic echelle modes. This includes 1) the derivation of a new absolute sensitivity for all primary settings, 2) a quantification of sensitivity variations as a function of time, 3) a characterization of the echelle blaze function (shape and shift) versus time and location on the detector (e.g., changes due to monthly offsettings), and 4) the determination of an on-orbit absolute sensitivity calibration for all secondary settings. I will quantify the improvements in terms of flux accuracies, and I will give an overview of the tools implemented into the pipeline in order to correctly take the above mentioned issues into account when calibrating STIS echelle data.

1. Introduction

The Space Telescope Imaging Spectrograph (STIS) has four echelle grating modes, namely E140H, E140M, E230H and E230M, which provide spectroscopic coverage from 1150 Å to 3100 Å at resolving powers from R $\sim$ 30,000 to R $\sim$ 110,000. Through simultaneous observations of multiple orders, these modes are designed to maximize the spectral coverage achieved in a single exposure of a point source.

The sensitivity calibration of the echelle modes has been particularly challenging over the years for several reasons. While an initial on-board absolute flux calibration for the 12 primary echelle medium- and high-dispersion modes is presented in Bohlin (1998), several important issues related to the echelle calibration have been tackled broadly and in depth only more recently. This has been possible thanks to a better characterization of the performance of the STIS instrument as a whole that has been performed in the meantime. Unless STIS is repaired, it is unfortunately too late for new STIS observations to benefit from such improvements. However, our deeper understanding of the flux calibration of the echelle modes will allow us to greatly enhance the scientific value of the echelle archival data, amounting to a total of about 30% of the STIS observations acquired during the whole lifetime of the instrument.

In this paper I will present the state of the art of the final sensitivity calibration for all the STIS spectroscopic echelle modes. I will only consider the most recent issues that have affected or are still affecting such calibration. For each issue, I will describe the effects on the spectra, the type of data that are in place in order to address them, the data analysis and results, the tools that have already been implemented or are going to be implemented into the calibration pipeline in order to apply appropriate corrections and what is still pending.

The four major outstanding issues for the sensitivity calibration of the echelle modes are an updated absolute flux calibration for all primary settings, a time dependent sensitivity, a characterization of the blaze function shift and shape as a function of time and location on the detector, and an on-orbit flux calibration for all secondary settings.

[1]On assignment from the Space Telescope Division of the European Space Agency

2. Updated Absolute Sensitivity Calibration for Primary Settings

The absolute sensitivity calibration of the STIS echelle spectroscopic modes is derived with an empirical method consisting in dividing the observed count rates by the known flux of a standard star. For all the echelle primary settings, this calibration is based upon observations of the pure hydrogen white dwarf G 191-B2B (STIS primary standard star), and the hot-subdwarfs BD +75°325 and BD +28°4211 (STIS secondary standard stars) and is detailed in Bohlin (1998). The observations were obtained by programs 7096, 7657, 7673, and 7917 during 1997 and 1998. The 0.2X0.2 aperture and the standard peak-up target acquisition were used for this purpose.

We have recently re-derived the sensitivity calibration of the echelle primary settings in light of the implementation of an improved algorithm for the subtraction of the scattered light. This new algorithm, developed by Lindler & Bowers (2001), turned out to be necessary in order to avoid unrealistic negative fluxes in the saturated cores of strong absorption lines, and was installed in the archive pipeline on December 21, 2000 with CALSTIS version 2.9 (Valenti et al. 2002). It is based on a two-dimensional scattered light model of the telescope and spectrograph, instead of being a simple one-dimensional linear interpolation of the minimum intensity between echelle orders. Its application reduces the corresponding flux uncertainties from about 10% to less than 2% (see Figure. 1 in Valenti et al. 2002).

Following the introduction of this 2D algorithm, the sensitivity calibration was re-determined in a self-consistent way. The new calibration differs by 2-7%, depending on the primary setting considered, compared to the calibration previously available. The changes, as expected, are not due to the improved scattered light subtraction, since heavily absorbed lines are not included in the sensitivity determination process. They are instead due to 1) an updated procedure that filters sharp absorption lines in the echelle spectra of the observed standard stars during the order-by-order sensitivity determination (see Bohlin 1998 for more details), 2) the removal of observations with non-zero monthly offsetting in order to avoid flux uncertainties related to the resulting shift of the blaze function, and 3) the inclusion of additional observations from the monitoring program 7673 spanning a longer time frame (up to middle 1999).

New PHOTTAB reference files were delivered to the pipeline on April 4, 2005. The time dependence of the sensitivity (see § 3) and the shift of the blaze function (see § 4) and were also backed out for self consistency. The new files reduce the flux uncertainty on the absolute flux calibration to less than a few %. Figure 1 shows the improvement in the flux calibration for observations of the standard star BD +28°4211 taken in 1997.7 with the grating E140M centered at 1425 Å.

3. Time-Dependent Sensitivity

A few representative primary settings of the four STIS echelle gratings were routinely monitored over the operational lifetime of the instrument in order to follow time dependent changes in sensitivity due to contaminants or other causes. The relatively bright ($V = 10.52$ mag) spectrophotometric standard star BD +28°4211 was used for this purpose with a 0.2X0.2 aperture. The observations were obtained twice a year as part of the Sensitivity Monitor calibration programs 7673, 8424, 8857, 8919, 9628, and 10033. The normal mode select mechanism (MSM) shifting was disabled for these monitoring observations in order to minimize variations due to spatial displacements of the spectra (see § 4). The echelle settings monitored in this way include E140H(1416), E140M(1425), E230H(2263), and E230M(1978, 2707).

The analysis of the monitoring observations for the low-dispersion modes already established clear evidence for wavelength dependent variations of the UV throughput over time (Stys, Bohlin, & Goudfrooij 2004 and references therein) at rates at times as high as

3% and 1-2% per year in the FUV and NUV, respectively. A comparison of the spectra of the same standard star taken at different epochs in the same setting indicates that a similar phenomenon affects echelle observations. Flux uncertainties as high as $\sim$ 15-20% are observed in the most recent datasets. Figure 2 shows an example of the flux of the standard star BD +28°4211 observed at different epochs in the E140M(1425) primary setting.

An analysis of the echelle monitoring data indicates that the primary settings investigated so far have a time-dependent sensitivity (TDS) which is similar to that observed for the first order modes (i.e., Figures. 6 and 7 in Stys et al. 2004). An example of these sensitivity changes over time is given in Fig. 3. A larger scatter, however, is observed in the echelle data compared to the first-order modes. We are still in the process of investigating the cause of such a scatter. This could be due, e.g., to 1) temperature effects in the FUV, as already observed in the first-order modes, 2) miscentering of the source in the smaller aperture used (0.2X0.2 versus the 52X2 slit of the first-order modes), or 3) a not properly corrected time dependence of the blaze shift (see § 4).

Since the monitored echelle primary settings behave similarly to the corresponding first-order modes in terms of TDS and a much larger number of observations is available for the latter to properly characterize the sensitivity change with time, we adopted the first-order mode correction for all the echelle settings. New TDSTAB reference files were delivered to the pipeline on April 4, 2005 in order to take this correction into account. In light of this delivery, the uncertainty associated to the absolute sensitivity of the echelle primary modes is reduced from about 15-20% to less than a few %, as shown in Fig. 4 for E140M(1425).

4. Blaze Function Calibration Issues

In echelle spectroscopic mode the spectrum of a target is broken up into several orders arranged as quasi-horizontal stripes on a two-dimensional detector (e.g., MAMA detectors). The characteristic efficiency along a single order is known as echelle blaze function or echelle ripple, because of its characteristic shape. Each spectral order needs to be corrected for the echelle blaze function before merging all the orders together to form a single calibrated spectrum. As already mentioned in § 2, this correction is performed by following an empirical approach in the case of STIS. Early observations of spectrophotometric standard stars are used to determine both the average sensitivity and the shape of the blaze function of each order by fitting a set of spline nodes to the ratio of the observed count rate divided by the standard star flux. However, there are additional complications for STIS. The echelle blaze function varies with location of the spectrum on the detector and with time since the instrument was installed on HST (e.g., due to contaminants or changes in the optics). These changes lead to saw-tooth wavelength dependent flux calibration errors that can be as large as about 15-20%. Figure 2 shows an example of how the scallopping of the spectra of the same standard star worsens with time due to blaze function variations not properly handled by the calibration pipeline.

Two types of changes can affect the blaze function, namely shifts and shape variations. Also, it is necessary to distinguish between effects due to location and effects due to time. This is the reason why a modular approach was selected where the spatial and temporal variations of the blaze function were tackled separately. Two different datasets were used for this purpose. A set of observations with the same zero monthly offsetting but different observational epochs was used to address temporal variations. A complementary set of observations taken in the same epoch with different monthly offsettings was instead analysed to determine spatial dependencies of the blaze function.

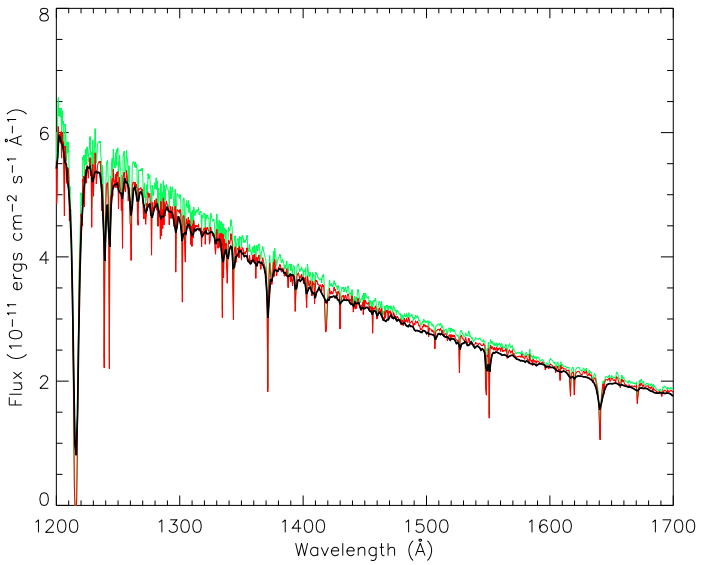

Figure 1: Flux of the standard star BD +28°4211 observed with HST/STIS in 1997.7 in the echelle primary setting E140M(1425) and calibrated with the old (green) and the new (red) PHOTTAB reference file. HST/FOS observations are also indicated (black). The new sensitivity calibration differs by about 6% and 3% at around 1250 and 1650 Å, respectively.

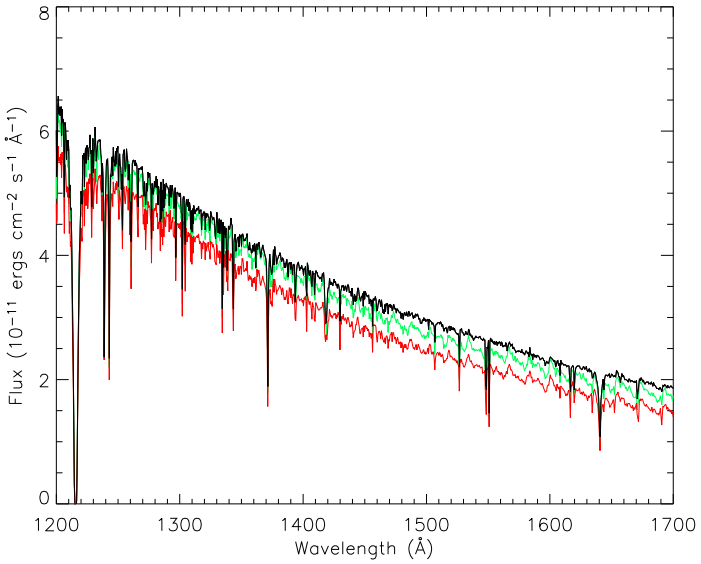

Figure 2: Flux of the standard star BD +28°4211 observed in the E140M(1425) echelle primary setting in 1997.7 (black), 2000.9 (green), and 2004.0 (red). The latest observations have a flux calibration which is about 4% and 15% lower at around 1250 and 1650 Å, respectively, than the observations acquired at the beginning of STIS operations. The scallopping of the spectra that worsens with time is due to an extrapolation within the pipeline of the blaze shift correction that does not properly handle datasets acquired after 2001 (see § 4 for more details).

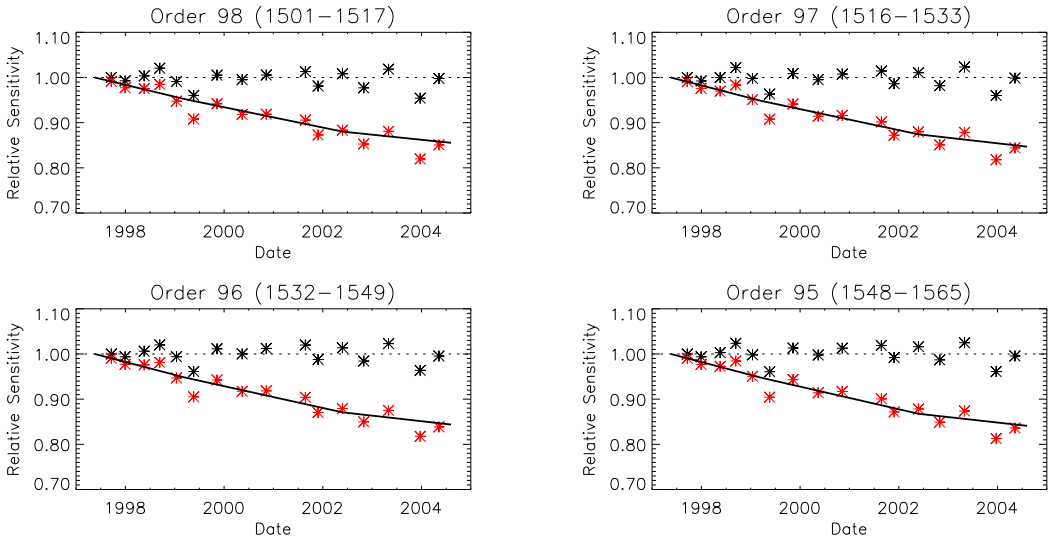

Figure 3: Sensitivity changes with time in selected orders of the E140M(1425) echelle primary mode covering the 1501 to 1565 Å wavelength interval (red asterisks). TDS-corrected sensitivities are also shown (black asterisks). The correction is performed by considering the piece-wise linear fit to the TDS of the low-resolution G140L grating (solid line).

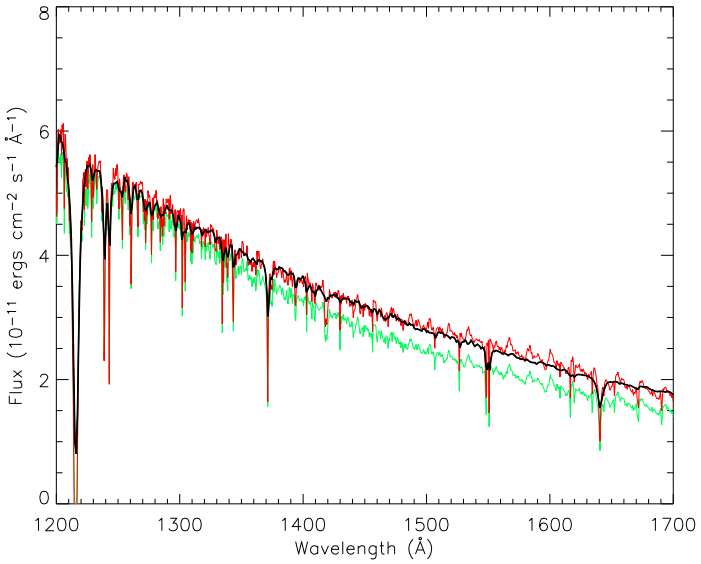

Figure 4: Flux of the standard star BD +28°4211 in the E140M(1425) echelle primary setting as observed in 2004 with (red) and without (green) TDS correction. The similarity of the echelle corrected flux compared to the HST/FOS flux (black) is remarkable. We are still in the process of properly correcting the scallopping of the spectra due to blaze shift effects.

4.1. Blaze Shift Corrections already implemented into the Pipeline

Bowers & Lindler (2003) already demonstrated with simple optics concepts that the blaze function shifts by a certain number of pixels when a spectrum is acquired with an offset of the STIS MSM that makes it fall onto a slightly different position on the MAMA detectors. In the past, monthly offsettings were used regularly for STIS observations in order to uniformly age the MAMA detectors. Also, the blaze function shifts by a different number of pixels compared to the position of the spectrum itself, i.e., compared to the wavelength scale.

An algorithm was also developed by Bowers & Lindler in order to predict a simple linear offset of the nominal calibration curve for each MSM shift. All archival non-proprietary echelle observations of sources with a continuum over the period from STIS launch to December 2001 were used for this purpose. The blaze shift of each dataset was inferred by shifting the echelle ripple pattern until the overlapping regions of adjacent orders are coincident. The shift was then correlated to the spectral shifts in both dispersion (Δx) and cross-dispersion direction (Δy) as determined from the accompanying spectral wave calibration exposures (wavecal). However, the 2-parameter linear model suggested by optics considerations and initially adopted by Bowers & Lindler, showed time-correlated residuals with an amplitude much larger than expected by simple TDS effects. As a consequence, a linear variation with time was also inserted in the algorithm and attributed to small changes in the grating itself.

The grating-dependent Bowers & Lindler's algorithm was implemented into the pipeline on September 5, 2002 with CALSTIS version 2.13b, and corrects all primary settings for the blaze shift. In addition, the monthly offsetting was turned off for echelle spectral modes starting in August 2002 in order to avoid the related flux uncertainties. While only earlier spectra are affected by the spatial component of the blaze shift, the temporal component has clearly continued to be present in all the echelle observations taken until the STIS failure in August 2004.

4.2. Temporal Variations of the Blaze Function

The time dependence of the blaze shift already implemented into the pipeline has been recently checked in order to verify that its extrapolation properly corrects echelle data acquired after 2001. The same monitoring observations of the standard star BD +28°4211 used to characterize TDS in five selected echelle primary modes (see § 3), were considered within this context. As previously mentioned, such observations were acquired with zero monthly offsetting over the whole period of STIS operations from 1997 to 2004.

For each setting and order, the blaze function was fitted with a third- or forth-order polynomial and a sigma-clipping algorithm to reject absorption lines, since the wavelength scale shifts differently compared to the blaze function. The fit was then normalized to its maximum in order to eliminate TDS and/or miscentering effects. Finally, the blaze shift of each order was obtained by cross-correlating the normalized fit of the blaze function at a certain epoch to the normalized fit of the blaze function in the first epoch available for each setting. The major result of this study is that the extrapolation of the time dependence of the blaze shift implemented into the pipeline does not hold anymore for reasons that differ from one grating to the other. The procedure we followed also allowed us to verify that the blaze function shape does not change with time within the flux uncertainties ($< 5\%$) of the standard star used.

In particular, an order dependence of the blaze shift was discovered for the first time. This dependence is mainly linear, even if second-order non-linear effects cannot be excluded, especially in the case of the E140H(1416) and E140M(1425) primary settings. Also, a different behavior of the blaze shift with time and order is observed depending on the grating considered. For E140H(1416), E140M(1425), and E230H(2263), the order dependence of the blaze shift gets stronger with time, with differences up to approximately 50 pixels between extreme orders in the worst case of the latest E140M(1425) dataset. An example of such a

trend is shown in Fig. 5 for the E230H(2263) primary setting. Since the Bowers & Lindler's algorithm is order independent, it gives a single blaze shift value which is an average over the orders and does not properly correct the extreme orders. Differences up to about 25 pixels between the pipeline value and the most appropriate blaze shift of such orders can be observed in the most recent observations. On the other hand, the order dependence of E230M at the central wavelengths 1978 and 2707 Å remains practically constant with time.

While E140H(1416), E140M(1425), and E230H(2263) have a blaze shift averaged over the orders that increases or decreases linearly with time, E230M centered at the two central wavelengths 1978 and 2707 Å, has an average blaze shift that clearly reverses its trend after the switch to Side 2 STIS operations in July 2001. This implies that the linear extrapolation of the Bowers & Lindler's algorithm is not suitable for data taken after that date. Errors as high as about 10 pixels are associated with such extrapolation. It is also interesting to notice that the two different primary settings of E230M have the same blaze shift behavior with time and order. This is an additional indication that the change with time is related to some physical parameter within the grating itself and not to wavelength dependent variations within the detector.

The implementation into the pipeline of a blaze shift time dependence that is also order dependent is still pending. New PHOTTAB reference file will be delivered in order to take such corrections into account. Once implemented, the new blaze shift correction will reduce the wavelength-dependent flux uncertainties to less than a few %. Figure 6 gives an idea of the type of improvements expected with this update.

4.3. Spatial Variations of the Blaze Function

The blaze function behavior versus the location on the detector was also investigated by using a homogeneous set of data from the calibration program 9619. The observations were simultaneously obtained with a 0.2X0.2 aperture and five different monthly offsettings in July 2002. Only the E230H(2513) primary setting was considered for this purpose. The data were analyzed with a procedure similar to the one used for the characterization of the temporal dependence of the blaze function. Once again, there is no indication of a change in the blaze function shape with position of the spectrum on the detector. Also, no strong dependence of the blaze shift on the order was found. Furthermore, the blaze shift averaged over the orders for a certain monthly offsetting, is similar to the value obtained with the Bowers and Lindler's algorithm of the pipeline, once the zero monthly offsetting is taken as reference, i.e., once the time component of the blaze shift is eliminated. This indicates that the blaze shift dependence on the detector location is already well characterized by the pipeline algorithm, at least for the E230H grating. The verification of the spatial component of the blaze shift adopted into the pipeline for the other gratings is still under review.

5. Calibration of the Secondary Settings

The last issue still pending for the sensitivity calibration of the STIS echelle modes is an improved calibration of the secondary settings. These settings are used almost as often as the primary settings. However, their absolute sensitivity calibration is still based on pre-launch data. Moreover, while TDS corrections have already been implemented into the pipeline for these settings based on the first-order mode trends, no attempt has been made so far to correct for blaze shifts. The similar behavior observed for the two different primary settings of E230M, seems to indicate that the blaze shift correction is in fact independent of the central wavelength, but we are still in the process of verifying this for the secondary settings.

On-board data of the primary spectrophotometric standard star G191-B2B within a 0.2X0.2 aperture were obtained for all primary and secondary settings in September 2001 as part of the calibration program 8915. The primary modes of such a homogeneous dataset

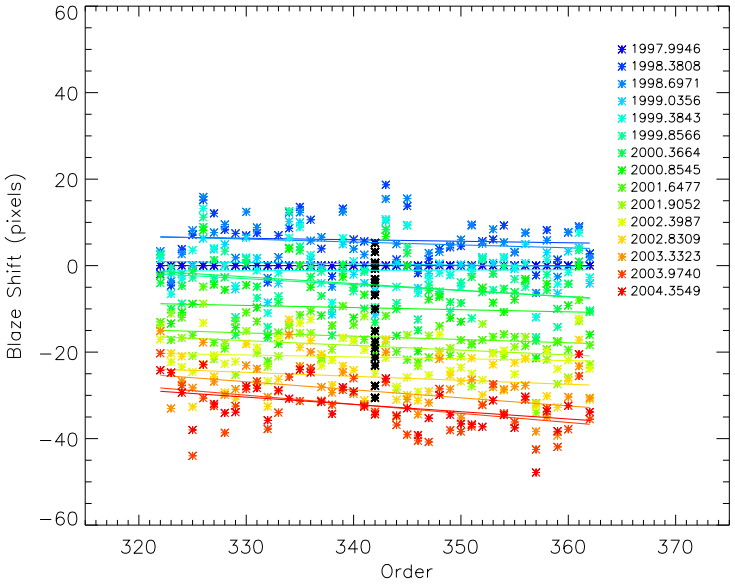

Figure 5: Blaze shift in pixels as a function of the order for every dataset analyzed in the E230H(2263) primary setting. The corresponding linear fit is also plotted for reference. See text for more details on how the shifts were derived. Different colors indicate datasets at different epochs. Blaze shift values associated to each dataset by the calibration pipeline are also indicated (black asterisks).

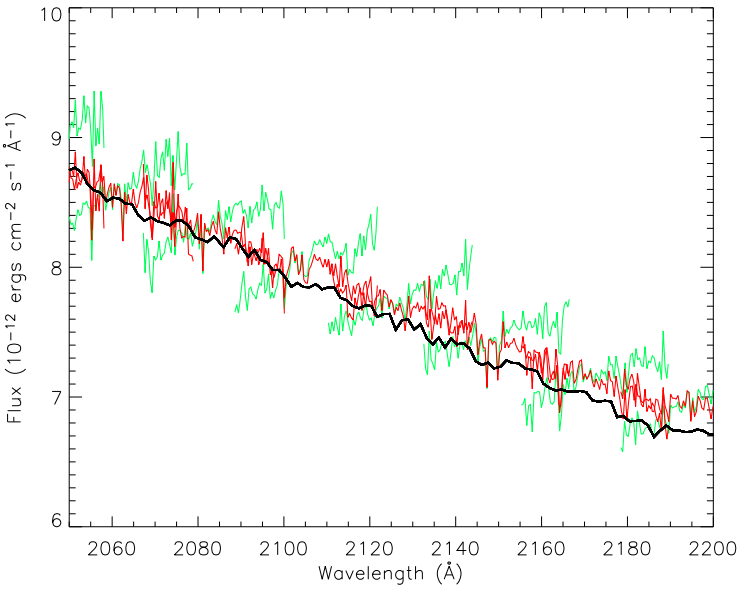

Figure 6: Flux of the standard star BD +28°4211 in the wavelength range 2050-2200 Å as observed in 2004.4 in the E230M(1978) primary setting with pipeline (green) and appropriate (red) correction of the blaze shift. The flux observed with HST/FOS is also indicated (black).

will be used to anchor all sorts of corrections of the secondary settings to the trends observed for the primary settings. The analysis of this dataset is still pending, but once completed it will give an on-board flux calibration for all echelle secondary settings that will reduce the associated errors from about 15% to 3-5%.

References

Bohlin, R. C. 1998, *Instrument Science Report* STIS 98-18 (Baltimore: STScI)

Bowers, C., & Lindler, D. 2003, in Proc. 2002 *HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 127

Lindler, D., & Bowers, C. 2001, BAAS, 197.1202

Stys, D. J., Bohlin, R. C., & Goudfrooij, P. 2004, *Instrument Science Report* STIS 2004-04 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

Valenti, J.A., Lindler, D., Bowers, C., Busko, I., & Kim Quijano, J. 2002, *Instrument Science Report* STIS 2002-001 (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Improvements to the STIS First Order Spectroscopic Point Source Flux Calibration

Charles R. Proffitt[1]

Science Programs, Computer Sciences Corporation, Baltimore, MD 21218

Abstract. During the operational lifetime of STIS, the primary emphasis of throughput measurements was the accurate determination of broad band point source fluxes observed using the low dispersion modes at the center of the 52X2 aperture, including measurement of time-dependent changes in sensitivity. We will review the current status and accuracy of this primary flux calibration, and discuss the remaining instrumental and model dependent uncertainties for both low and medium dispersion first order modes.

In practice, the bulk of STIS 1st order spectroscopic observations used apertures smaller than the 52X2, and in recent years a substantial fraction of observations were done at the E1 aperture positions, which were placed closer to the readout to minimize CTI losses. Analysis has now shown that *additional grating dependent throughput corrections are needed for these smaller apertures*, and vignetting corrections are needed for observations done at non-central positions on the detector, including those observations done at the E1 positions. These corrections can amount to several percent, and failure to include them can result in flux inconsistencies between observations done with different gratings or apertures. Once these corrections are applied, the flux calibration of well centered observations done using the commonly used 52X0.2 aperture should agree with the primary wide aperture calibration to better than 2%.

1. Overview of Primary Flux Calibration

The absolute flux calibration of Space Telescope Imaging Spectrograph (STIS) first order spectroscopic modes was discussed in detail by Bohlin (2000) and is primarily based on observations of the DA white dwarf (WD) standard stars G 191-B2B, HZ 43, GD 153, and GD 71 (Bohlin, Colina, & Finley 1995). The fundamental assumption of these calibrations is that the spectral energy distribution of the nearly pure hydrogen atmospheres of these stars can be calculated with great precision (Barstow et al. 2001), allowing them to be used as absolute flux standards.

In addition to the observations defining the primary sensitivity calibration, routine monitoring was also done to follow time dependent changes. Rather than using one or more of the fundamental standards, two other stars, (GRW+70°5824 for the STIS G140L and G230L modes, and AGK+81°266 for the other first order modes), were selected that give good signal-to-noise with short exposures and which can be observed at any time of the year. These monitor observations established clear evidence for wavelength dependent variations of the optical throughput over time (Stys et al. 2002 and references therein) at rates as high as 3% per year, as well as the need to correct CCD spectroscopic observations for charge transfer inefficiency losses (CTI) (Goudfrooij & Kimble 2002; Bohlin & Goudfrooij 2003).

[1]also Space Telescope Science Institute, and Catholic University of America

The original CALSTIS software and reference files have been modified to correct for these changes.

Stys, Bohlin, & Goudfrooij (2004) reported on time-dependent sensitivity (TDS) changes measured through late 2004. The most notable change from previous trends was a distinct flattening that occurred beginning in early 2002 of rate of the sensitivity decline at most UV wavelengths. An example of these sensitivity changes over time is shown in Figure 1. Updates to CTI corrections for STIS spectroscopic observations are discussed elsewhere in this volume (Goudfrooij & Bohlin 2005).

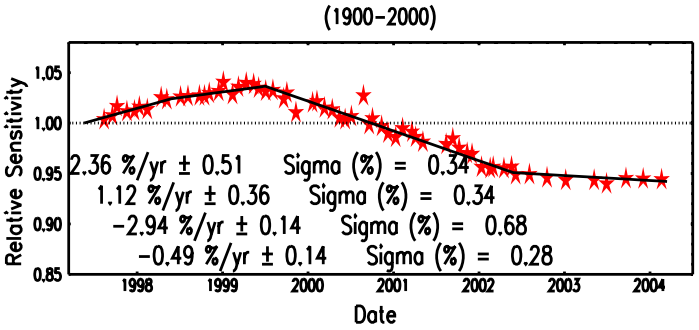

Figure 1: In this figure from Stys et al. (2004), the measured changes of the sensitivity with time are shown for the 2000 to 2100 Å wavelength interval of the G230L grating, along with the best piece-wise linear fit.

For well exposed observations taken at the center position with the 52X2 aperture[2] using STIS low dispersion modes, once TDS and CTI corrections and a small temperature dependence (see Stys et al. 2004) of the sensitivity are taken into account, fluxes as measured over 50 - 100 Å wavelength bands repeat to better than 1%. However, once the uncertainties in the stellar parameters and stellar atmosphere models (Barstow et al. 2001) are taken into account, the absolute accuracy of the STIS first order low dispersion flux calibration for observations using the 52X2 aperture is estimated to be about 2% in the visual, and 4% in the far-UV (Bohlin, Dickinson, & Calzetti 2001). Because the primary limitations on the flux accuracy are systematic uncertainties in the white dwarf models, relative STIS fluxes and derived colors can be much more precise (e.g., see Maíz-Apellániz 2005).

1.1. Primary Flux Calibration of Medium Resolution First Order Modes

Because of the much larger number of central wavelength settings, and the presumption that the low dispersion modes would be preferred when highly accurate flux measurements were needed, the flux calibration of the medium resolution first order modes has been a lower priority than that of the low dispersion modes.

For the G140M and G230M gratings, observations of the standard DA white dwarf GD 71 obtained in 1997 and 1999 were used for the initial calibration. A few observations of another white dwarf standard GD 153 were also obtained in 1997. Subsequently the primary white dwarf standard G 191-B2B was observed in 2000 and 2001. We recently reconsidered all of this data in deriving new photometric conversion table (PHT reference file) throughput curves for all MAMA medium resolution first order modes.

For the G230MB, G430M, and G750M gratings, the initial on-orbit calibration was done using observations of the hot-subdwarf BD+75°325 obtained by programs 7094, 7656, and 7810 during 1997 and 1998. This star is not one of the primary white dwarf standards,

[2]STIS clear apertures names give the size of the aperture in arc-seconds, with the first number giving the dimension in the cross dispersion direction, and the second the size in the dispersion direction.

but is instead a secondary standard which is tied to the primary standards through low-dispersion STIS observations. This procedure results in a reference spectrum that is of lower resolution than the medium resolution modes it is used to calibrate, and also ties the medium resolution calibration to the low resolution one. We recently rederived new pht curves for the CCD medium resolution first order modes using spectra of G 191-B2B obtained by programs 8421 and 8916 between January 2000 and April 2002.

For the G140M grating settings, the setting of the STIS mode select mechanism (MSM) which holds the grating was changed a number of times. During pre-launch testing, the MSM had been set to place G140L and G140M spectra near row 530, but this was too close to the shadow of the FUV-MAMA repeller wire and in March 1997, shortly after STIS's installation into HST, the grating tilt was changed to put the spectra near row 600 of the detector. In January 1998 an additional small monthly offsetting was added for all MAMA spectroscopic modes to prevent over use of any single portion of the detector. However, the glow that increases the dark current for the FUV-MAMA detector was strongest in the upper part of the detector, and so in March 1999, the MSM grating tilts for the G140M and G140L detectors were changed again, this time to place 1st order spectra near row 400 – below the repeller wire and in a part of the detector with significantly lower dark current. Unfortunately, at most central wavelengths, GD 71 has been observed only near row 600 and G 191-B2B only near row 400, making it difficult to disentangle vignetting differences at the two positions from differential errors in the model stellar spectra or from inadequacies in the treatment of time dependent sensitivity changes in STIS. This is complicated by the significantly lower signal-to-noise and the larger number of unmodeled absorption lines in the G 191-B2B spectra.

The region near Lyman-alpha presents special difficulties, as this is the region where the detailed WD model spectra are most likely to be in error. The ratio of the G140M observations of GD 71 and G 191-B2B at the 1222 setting to the model spectra cannot be made consistent with the same sensitivity function. The modeled Lyman-alpha line in GD 71 is a too wide relative to that of G 191-B2B to be consistent with the observations. While this may in part be due to the different MSM settings used for the observations of these two stars, the model spectra do not perfectly reproduce the actual line profile. To avoid these uncertainties, all wavelengths between 1205 and 1240 angstroms were excluded from the fit and only 3 spline nodes were used in the fit. This smoothly interpolates the sensitivity curve over the region of the Lyman-alpha line, but may result in relative flux errors of order 5%.

Similar, but more modest caveats apply to some some G430M settings where the G 191-B2B model spectra contain significant absorption lines. In these cases, retaining most wavelengths affected by the lines, results in a reasonable fit, but changes to the model spectrum could affect the fit at the 1 to 2% level.

G750M spectra at wavelengths larger than 7000 angstroms are affected by fringing in the CCD substrate. Because of the large tilt of G750M spectra, the STScI provided tasks for defringing G750M spectra were only designed to work with 2D rectified files and are not applicable for 1D spectral extractions. To defringe these data we performed 1D spectral extractions of tungsten lamp spectra done with the 0.3X0.09 aperture, and then divided out a 3 node spline fit to remove any lamp vignetting from the fringe flat. This normalized fringe spectrum was then divided into the corresponding WD spectrum. While not perfect, this provides adequate fringe removal for our purposes. The new throughput curves for G750M tend to have fewer wiggles than do the previous calibration. This may be due to the inclusion of better fringing corrections in our analysis.

For a few rarely used central wavelengths, no observations of primary standard WD stars are available to determine the throughputs. These include G140M 1218 (17 external observations), and 1400 (7 obs.), G430M 4781 (2 obs.), and G750M 10363 (20 obs.), as well as a few central wavelengths which were never used for external targets (G140M 1387, 1540, 1640; G230M 2600, 2800, 2828; G750M 9286, 9806).

At most wavelengths, the changes from the previous calibration are modest.($< 3\%$), and the overall flux accuracy should be comparable to that for low dispersion modes. The G140M calibration is slightly more problematic due to the multiple MSM tilts and uncertainties in the modeling of the Lyman-alpha profile in WD atmospheres. These uncertainties may cause additional localized flux errors of 2 to 5%.

2. Aperture and Position Dependent Corrections

In practice, many STIS observations were not taken at the central position using the wide photometric 52X2 aperture, and the absolute flux calibration for such observations has not been as intensively studied. Below we will consider the corrections that are needed for other apertures and for other positions on the detector.

2.1. Vignetting and Low-order Flat Fields

STIS first order modes were usually used with long slits that span the CCD or MAMA detector. For some observations, multiple targets were placed at different locations along the slit, and for the CCD detectors, new E1 aperture positions were defined closer to the readout register, near row 900, to minimize CTI losses. The throughput as a function of wavelength for different positions along the aperture can in principle vary either because of changes in the optical vignetting, or because the sensitivity of the CCD or MAMA detector at a given wavelength varies over its surface.

To measure such effects, a number of calibration programs were done in which a bright star was dithered up and down the length of the 52X2 aperture. In addition, for the CCD modes, sensitivity monitor measurements made after April 2002 routinely included measurements at both the central and E1 aperture positions. The ratio of E1 sensitivity to central row sensitivity is therefore very well measured, while at other positions along the slit, more limited data is available, with 1 to 4" spacing along the cross dispersion direction.

These observations were used to derive low order flat fields for the G140L, G230LB, G430L, and G750L gratings. The medium resolution gratings lack sufficient data at most central wavelengths, while the G230L mode does not appear to require a low order flat field correction. The low order flat fields are defined to be unity along the same row where the sensitivity function calibration was set. In Figure 2 we compare the ratio of observations made at the 52X2E1 and 52X2 aperture positions with the same ratios from our adopted low order flat field images.

These low order flat fields do have some limitations, and are not intended to follow every small scale wiggle in the measured sensitivity ratios. At locations other than the central and E1 positions only limited and coarsely spaced measurements of the sensitivity variations are available. Also, near the short wavelength end of both the G430L and G750L, structures in the throughput curves from interference in the coatings on the order sorter filters on these gratings make the flux calibration extremely sensitive to the exact target positioning. The low-order flat fields are strictly applicable to targets positioned along the center line of the long apertures. Targets that are significantly offset in the dispersion direction will not be corrected with the same degree of accuracy.

For first order medium resolution modes, vignetting data is only available for a limited number of central wavelengths and has not yet been fully analyzed. For these modes no low order flat field corrections are currently supplied in the pipeline.

2.2. Corrections that Depend on the Grating-Aperture Combination.

Adopted aperture throughputs as a function of wavelength were determined by Bohlin & Hartig (1998) using a model of the STIS PSF and assuming that the throughput of a given aperture is only a function of wavelength. Bohlin & Hartig noted that small aperture

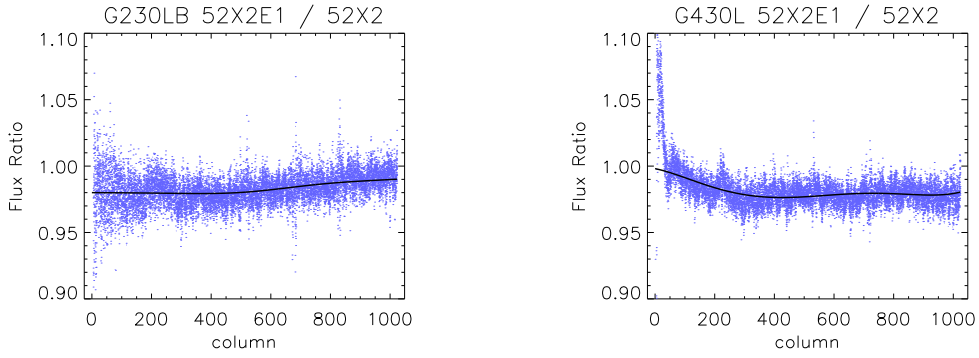

Figure 2: If no lflat corrections are applied, fluxes measured at the 52X2E1 aperture position are systematically low compared to those measured at the standard 52X2 position. Here we show for the G230LB and G430L the ratio at each wavelength for each of the 14 sensitivity monitor visits which measured the sensitivity at both aperture positions (blue dots). The solid lines show the value of the adopted lflats for each grating at the E1 position.

throughputs for the G430L and G750L did seem to be systematically high by a few percent, and speculated that this might be related to the Lyot stops that are attached to these gratings. However, since the smaller apertures also show significant throughput variations due to telescope breathing and small offsets from the center of the aperture, further investigation of this effect was deferred. It was assumed that observations needed high absolute flux accuracy would use the 52X2 aperture. However, subsequent experience has shown that the narrower 0.2" wide aperture positions have been used about as often as the 2" ones.

Because of the heavy use of the 52X0.2 and 52X0.2E1 aperture positions, we have carefully reconsidered how the aperture throughputs vary for each grating. We identified a number of well centered calibration observations where in a single visit the same grating was used on the same target with both wide and narrow apertures using the same part of the detector. All targets and peakups for a given set of observations to be compared were done either at the center position *or* at the E1 position, and all peakups were done using 0.1" or smaller apertures. The results of comparing the measured 52X0.5 and 52X0.2 fluxes to the 52X2 aperture fluxes are shown in Figure 4. If the residuals had smoothly varied as a function of wavelength, they might have been explained as an error in the adopted aperture throughput as a function of wavelength. Instead we see inconsistencies of up to 6% in the overlap region between the different CCD gratings – the effective aperture throughput does depend on which grating is being used. In addition, the 52X0.2/52X2 throughput ratio is about 2% smaller when the comparison is done at the E1 rather than the central position. This may be due to small variations in the width of the 52X0.2 slit along its length.

Similar comparisons for the smaller 52X0.1 and 52X0.05 apertures are shown in Figure 5. While there is some evidence for the same offsets between gratings seen in the 52X0.2 aperture throughput, the exposure-to-exposure variations caused by telescope breathing and minor mis-centering is often larger than this effect. The throughput at long wavelengths is also systematically smaller than predicted by the nominal throughput curves.

Rather than installing an aperture-dependent grating throughput or a grating-dependent aperture throughput in the current calibration architecture, it was decided that it would be simpler to define a new flux correction vector that depends on both the aperture, grating, and central wavelength being used. Such a correction vector can also be used to correct for the small scale sensitivity differences between the central and E1 position that were not put into the low order flat field. Figure 6 shows examples of this for the G430L grating.

In Figures 7 and 8 we show extracted fluxes for G230LB, G430L, and G750L observations of the A star HD 141653 (HR 5886) taken using the 52X0.2 aperture. Applying

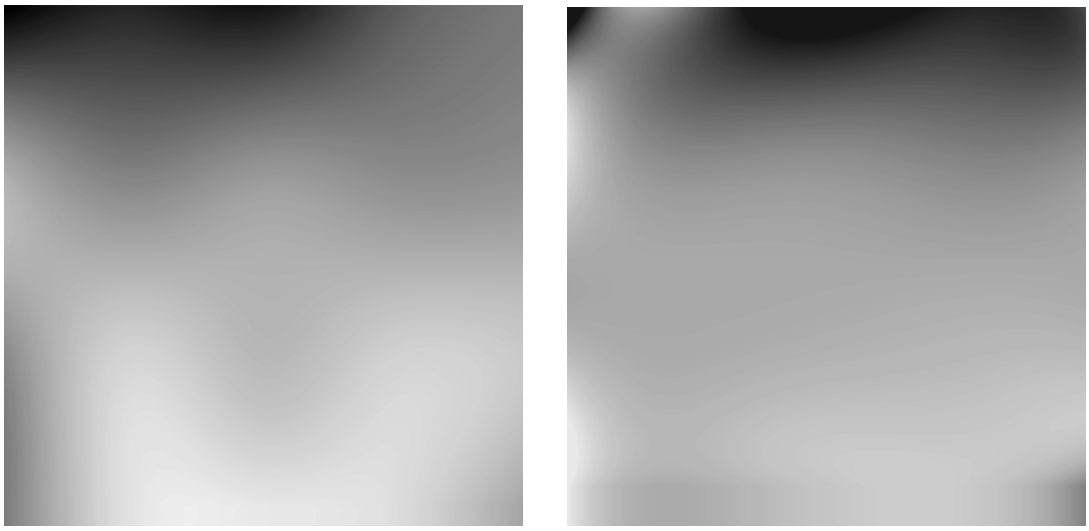

Figure 3: Adopted low order flat field images for the G230LB (left) and the G430L (right) are illustrated. The value of the G230LB lflat ranges from 0.964 to 1.012, the G430L from 0.936 to 1.019, and the G750L (not shown) from 0.864 to 1.018.

the new correction significantly improves the agreement in the overlap region between the spectra. Note that we do not attempt to correct for the short wavelength fringing due to the order sorter filters on the G430L and G750L. As a result, the first $\sim$15 pixels of the G430L and the first $\sim$25 pixels of the G750L are not as accurately corrected.

The throughput of 0.2" and smaller apertures at longer wavelengths is also especially sensitive to minor changes in the telescope pointing and breathing, and it is easy to find examples where the application of the new correction makes the agreement in the overlap region between the G430L and G750L worse rather than better.

While the throughput discontinuities between gratings cannot be explained by a simple error in the aperture throughput as a function of wavelength, the need to determine the GACTAB corrections empirically makes it difficult to completely disentangle modest errors in the small aperture throughputs from the effects of the Lyot stops. Since the medium resolution CCD gratings lack the Lyot stops found in the G430L and G750L, they should not need these corrections, and could in principle be used to more directly check the small aperture throughputs. However, only a handful of observations are available where an medium resolution spectrum can be directly compared to a spectrum taken with a larger aperture or with a low dispersion grating using the same aperture. These observations suggest that errors in the 52X0.2 aperture throughput are less than 2% over most of the relevant wavelength range, but may be larger at wavelengths longer than 7000 Å.

It might be expected that any vignetting due to the G430L and G750L Lyot stops would behave differently for extended and point-source targets, since for extended targets the far PSF wing from source flux outside the aperture may compensate for much of the flux from the part of the target "in aperture" that is lost to the PSF wing. There are few data sets that can be used to test extended sources for narrow aperture flux inconsistencies, but a good test case is provided by a set of long slit observations of Saturn and its rings which used the G230LB with the 52x0.5 aperture and the G430L and G750L with the 52X0.1. A comparison of the summed x2d fluxes shows no evidence of the flux inconsistencies seen for small aperture point source observations. The throughput corrections applied for the effects discussed above will therefore be implemented only as part of the x1d flux extraction, but not for the fluxes derived as part of x2d processing.

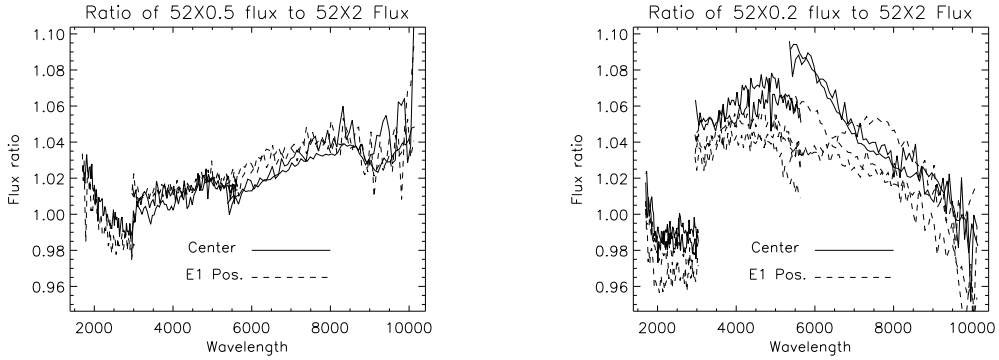

Figure 4: Ratio of measured narrow to wide 52X2 aperture fluxes after correcting for the nominal aperture throughput vs wavelength. Results are show for the 52X0.5 and the 52X0.2 apertures. Dashed lines show the ratio as measured at the E1 aperture position, while solid lines show the ratio at the central position.

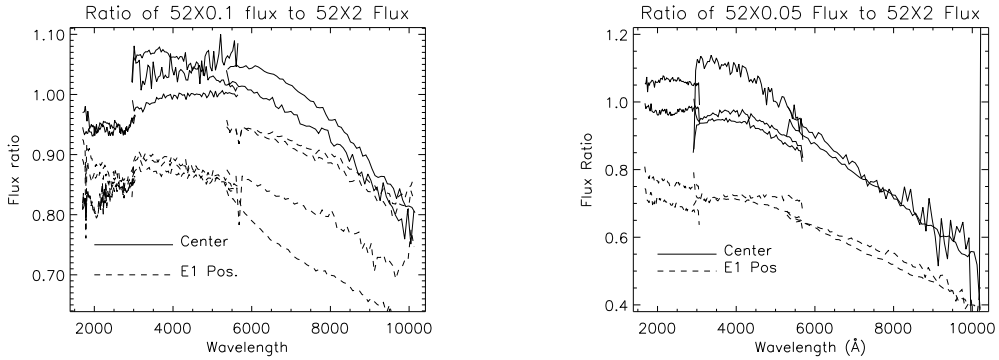

Figure 5: Same as Fig. 4, except for the 52X0.1 and 52X0.05 apertures.

2.3. Effects of Being Offset from the Aperture Center

An analysis of peakup observations done at the 52X0.1E1 and 52X0.05E1 aperture positions showed that the E1 positions had originally been defined $\sim 2/3$ of a pixel off the aperture centers. This error was not corrected in the Science Instrument Aperture (SIAF) file until September 2003. Any observations at the E1 aperture positions prepared for flight before this update would have this mis-centering, unless the science observation was proceeded by a small aperture peakup at the E1 position. To test the effect that this has on the flux throughput, observations of the star BD+75°325 at the 52X0.2E1 aperture position with the G230LB, G430L, and G750L gratings were taken with both good centering after a peakup in a smaller aperture, and with a deliberate 0.033" offset (see Fig. 9).

While it might be initially surprising that at some wavelengths the throughput is higher for off-center observations than for well centered ones, this is simply explained by considering how the first Airy ring of the PSF illuminates the aperture. In Figure 9 we also show the predictions of a simple model, where a series of monochromatic point spread function images created using the Tiny Tim PSF modeling software (Krist & Hook 2004) are truncated by a mask matching the size of the aperture. Near 4500 Å, the first Airy ring just fits inside the 0.2" wide aperture, and any small miscentering results in substantial throughput loss. Near 8000 Å the sides of the first Airy ring are just outside the aperture when the target is well centered, and a small offset picks up more flux from the part of the Airy ring moving into

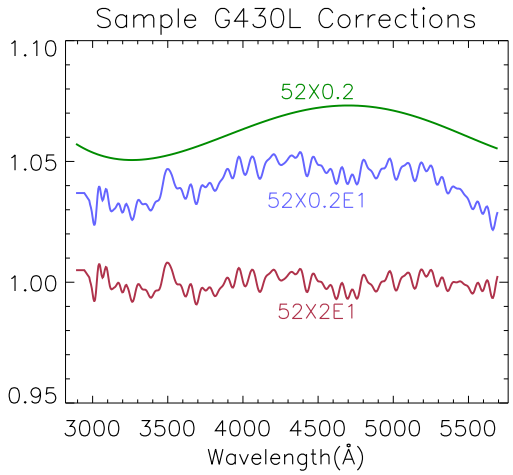

Figure 6: Examples of GACTAB correction vectors. Note that the small scale structure on the two E1 apertures is very similar, as this structure is intended to correct for the small scales differences between the E1 and regular position that remain after the vignetting corrections are applied. For the regular aperture positions, the GACTAB is only intended to correct for low-order differences in throughput for particular aperture/grating combinations

the aperture than is lost from the other side. The simple model is in qualitative agreement with the observations, although the observed effect is larger than the predicted one.

For off-center observations, the throughput is expected to be especially sensitive to small shifts or breathing changes, and so we make no attempt to correct for the effects of aperture miscentering in the pipeline.

3. Implementation of New Corrections in CALSTIS and OTFR

A new reference file type (GACTAB) was defined to handle the throughput corrections discussed in section 2.2 that depend on the particular combination of grating and aperture. Currently, we have only defined these corrections for the CCD 1st order modes and have only delivered corrections for the long-slit low dispersion CCD modes. Rows in this reference file are selected by the combination of optical element (grating), aperture, and central wavelength setting and contain a throughput correction vector as a function of wavelength which is applied in the x1d extraction during the conversion of net counts to physical flux.

Corrections for the temperature dependence of the throughput have been added to the time-dependent sensitivity (TDS) reference file, and these new corrections will be applied to the x1d extracted fluxes along with the existing TDS corrections.

The changes to the CALSTIS code needed to support these new and modified reference files are included in version 2.19 of CALSTIS which was released as part of STSDAS 3.4 in November 2005. This version of CALSTIS was installed in the OTFR pipeline as part of the OPUS 2006.1 release in December, 2005. This release also allows the GACTAB reference file to be made available to OTFR and puts the name of that reference file in the GACTAB keyword of the science image headers so that it can be used by CALSTIS.

Corrections for the low-order vignetting discussed in section 2.1 are done by including low order flat fields (LFL reference files) that are applied as part of the flat fielding correction. At the E1 positions, the LFL file will increase the final extracted flux values, while for narrow apertures the GAC corrections will decrease the flux values. Since these corrections work in opposition directions at the 52X0.2E1 aperture position, the delivery of the new CCD LFL files was timed to coincide with the implementation of the GAC correction.

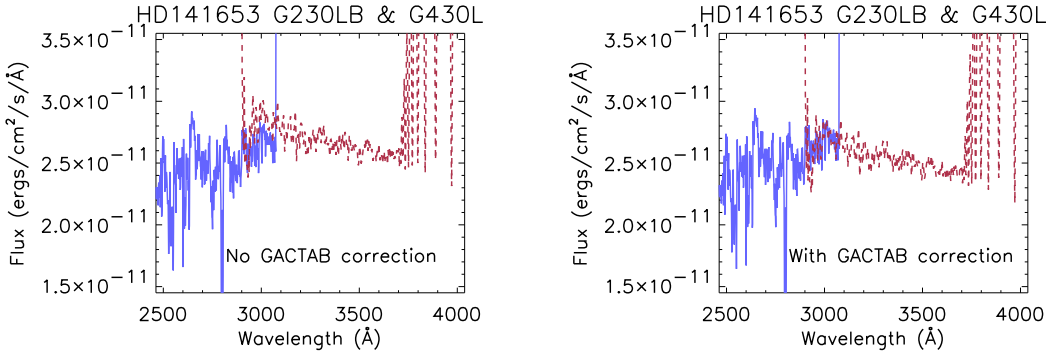

Figure 7: The extracted flux in the overlapping region of G230LB (solid line) and G430L (dashed line) STIS spectra of HD 141653 taken with the 52X0.2 aperture are compared without (left) and with (right) the inclusion of the new GACTAB correction vector for these grating/aperture combinations.

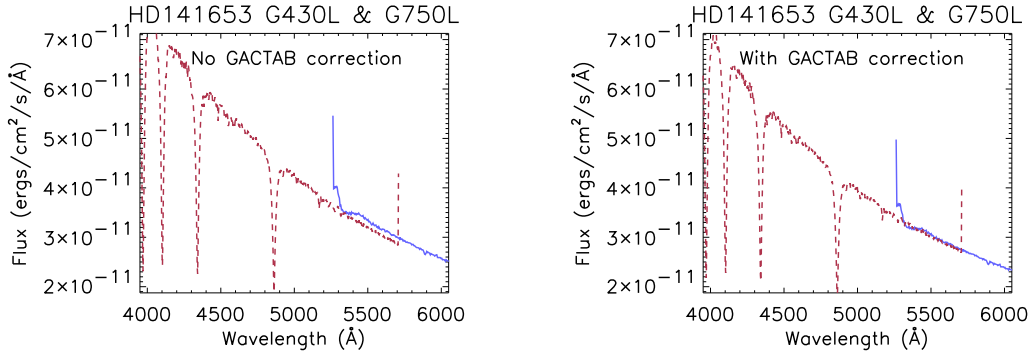

Figure 8: The same as Fig. 7, but for the overlap of the STIS G430L (dashed line) and G750L (solid line) gratings.

4. Remaining Issues for the Flux Calibration of STIS First Order Spectra

Final corrections for CTI effects are still pending (see Goudfrooij & Bohlin 2005). Once these are resolved then a final determination of CCD sensitivity and TDS corrections will be done using all available sensitivity monitor data up through the failure of STIS on 3 Aug 2004. An effort is also underway (Finley 2005) to improve the models used for the primary WD flux calibrations to allow absolute flux determinations of 1% or better at all wavelengths. Should these improved models be available before the final calibration of STIS data is undertaken, they will also be used in the final sensitivity determination.

Small aperture throughputs for MAMA first order modes as well as for the limited available data for medium resolution first order modes are still being reviewed. This may lead to some modest revisions of the base throughput curves for small apertures, and corresponding changes to the GACTAB corrections.

The flux calibration for well centered STIS first order spectroscopic observations done using the 52X0.2 aperture should be within 2% of the accuracy obtainable using the wider 52X2 aperture. Unfortunately, for observations where the centering is uncertain, including most observations done using the 52X0.2E1 aperture that were not preceded by a peakup in a smaller aperture, additional errors in the absolute flux of up to 5% may be encountered.

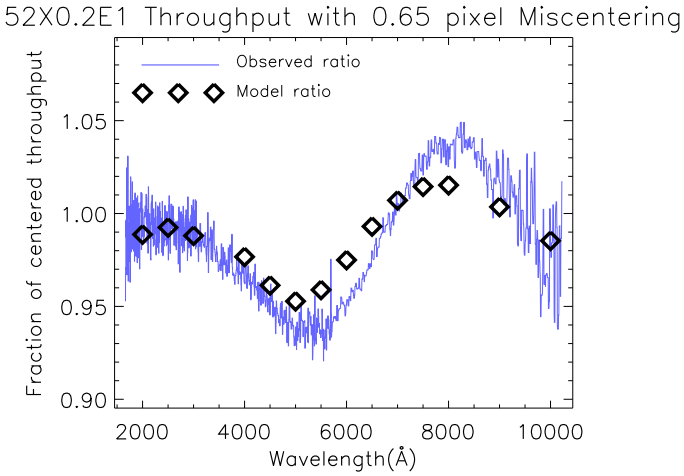

Figure 9: The effect of being 0.65 CCD pixels off-center in the 52X0.2E1 aperture on the flux throughput was measured by comparing well centered and mis-centered G230LB, G430L, and G750L spectra of the standard star BD+75°325. The diamond symbols show the prediction a simple model for the expected flux ratio.

However, as discussed by Maíz Apellániz (2005), even for these observations, broadband relative fluxes may still be highly accurate.

References

Barstow, M. A., Holberg, J. B., Hubeny, I., Good, S. A., Levan, A. J., & Meru, F. 2001, MNRAS, 328, 211

Bohlin, R. C. 2000, AJ, 120, 437

Bohlin, R. C. 2002, 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 115

Bohlin, R. C., Colina, L., & Finley, D. S. 1995, AJ, 110, 1316

Bohlin, R. C., Dickinson, M. E., & Calzetti, D. 2001, AJ, 122, 2118

Bohlin, R. & Goudfrooij, P. 2003 *Instrument Science Report* STIS 2003-03R (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

Bohlin, R. & Hartig, G. 1998, *Instrument Science Report* STIS 98-20 (Baltimore: STScI)

Finley, D. 2005, HST Proposal (ID#10654), 6974

Goudfrooij, P., & Bohlin, R. C.2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 181

Goudfrooij, P., & Kimble, R. A. 2002, 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 105

Krist, J., & Hook, R. 2004, *The Tiny Tim User's Guide* (Baltimore:STScI)

Maíz Apellániz, J. 2005, ArXiv Astrophysics e-prints, arXiv:astro-ph/0510785

Stys, D. J., Walborn, N. R., Busko, I., Goudfrooij, P., Proffitt, C., & Sahu, K. 2002 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 205

Stys, D. J., Bohlin, R. C., & Goudfrooij, P. 2004, *Instrument Science Report* STIS 2004-04 (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The HST/STIS Next Generation Spectral Library

M. D. Gregg[1]

Department of Physics, University of California, Davis, CA 95616

D. Silva

European Southern Observatory, Karl-Schwarzschild-Str. 2, D-85748 Garching, Germany

J. Rayner

Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822

G. Worthey

Department of Physics, Webster Hall 1245, Washington State University, Pullman, WA 99164

F. Valdes

CCS, National Optical Astronomy Observatory, P.O. Box 26732, Tucson, AZ 85726

A. Pickles

Caltech, 1200 E. California Blvd., 105-24, Pasadena, CA 91125

J. Rose, B. Carney

Dept. of Physics and Astronomy, University of North Carolina, CB 3255, Chapel Hill, NC 27599

W. Vacca

SOFIA-USRA, NASA Ames Research Center, MS N211-3, Moffett Field, CA 94035-1000

Abstract. During Cycles 10, 12, and 13, we obtained STIS G230LB, G430L, and G750L spectra of 378 bright stars covering a wide range in abundance, effective temperature, and luminosity. This HST/STIS Next Generation Spectral Library was scheduled to reach its goal of 600 targets by the end of Cycle 13 when STIS came to an untimely end. Even at 2/3 complete, the library significantly improves the sampling of stellar atmosphere parameter space compared to most other spectral libraries by including the near-UV and significant numbers of metal poor and super-solar abundance stars. Numerous calibration challenges have been encountered, some expected, some not; these arise from the use of the E1 aperture location, non-standard wavelength calibration, and, most significantly, the serious contamination of the near-UV spectra by red light. Maximizing the utility of the library

[1]Participating Guest, IGPP, Lawrence Livermore National Laboratory

depends directly on overcoming or at least minimizing these problems, especially correcting the UV spectra.

1. Introduction

A wide variety of astrophysical problems can be addressed only through spectral modeling (see Leitherer et al. 1996), from constraining the turnoff age of ellipticals to deriving the initial mass function in starbursts. In almost all galaxy spectral modeling applications, the limiting reagent is the input spectral library because none of the extant spectral compilations cover luminosity, metallicity, and T_{eff} space sufficiently, especially at UV wavelengths.

To help fill the need for a better stellar library, we have carried out a STIS snapshot program (GO9088, GO9786, GO10222) to construct a "Next Generation" Spectral Library (NGSL). Stellar targets were selected from various metallicity catalogs, and nearly all have $V < 12.$ and HIPPARCOS parallaxes. The observations used the low dispersion gratings and CCD detector of STIS to take high S/N spectra of stellar targets distributed over four metallicity groups:

very low	$[\mathrm{Fe/H}] < -1.5$
low	$-1.5 < [\mathrm{Fe/H}] < -0.5$
near-solar	$-0.3 < [\mathrm{Fe/H}] < +0.1$
super-solar	$+0.2 < [\mathrm{Fe/H}]$

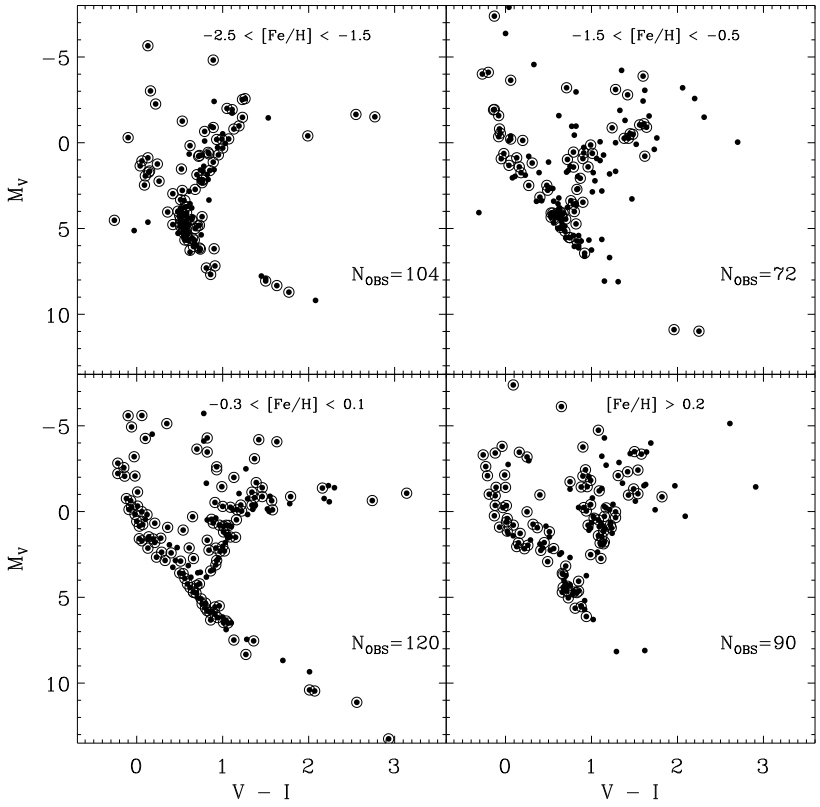

Figure 1: The final NGSL sample: circled points are stars for which we obtained STIS spectra; uncircled are the stars denied fame by the failure of STIS. The super-solar and low metallicity samples are somewhat under-represented because only solar and extremely metal poor stars were done in the first cycle of observing.

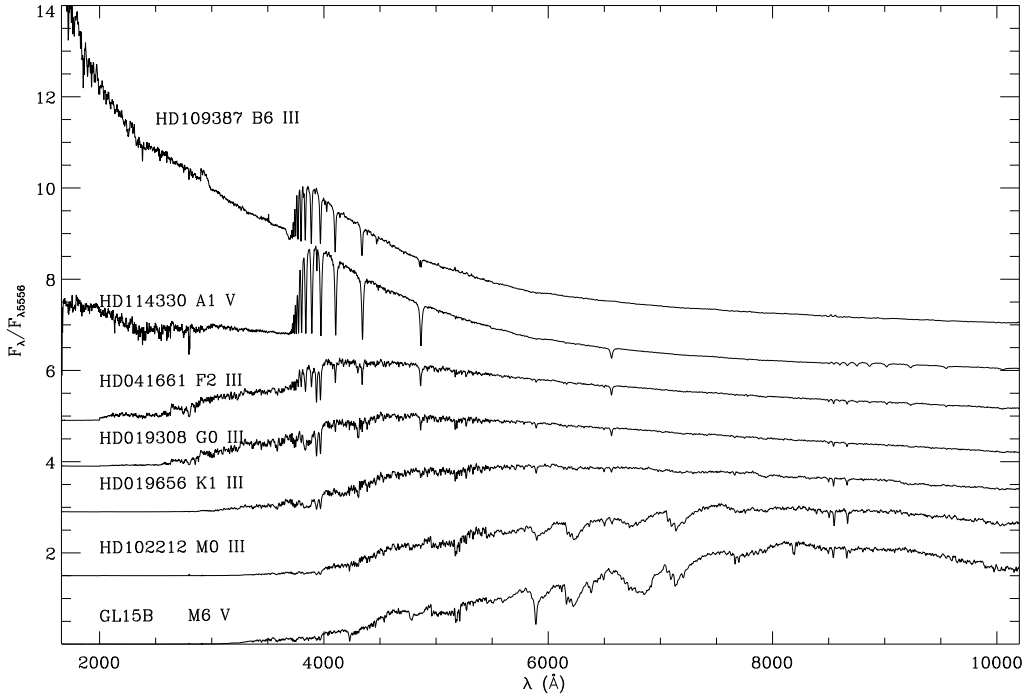

Figure 2: Example flux calibrated STIS NGSL spectra showing the wide range of spectral types covered by the program stars.

The three low dispersion CCD modes of STIS, G230LB, G430L, and G750L, cover the wavelength range 1670-10250Å. The original goal was to obtain 600 stellar spectra covering $T_{eff} - L - Z$ space far better than has been done before. This would have been reached in Cycle 13, but the failure of STIS in August, 2004, halted data collection at 380 stars. Even with the reduced number, the targets sample most of the HR-diagram in each bin (Figures 1, 2). Figure 3 compares two early K giants with different metallicites, demonstrating the need for a wide range of abundance in a stellar library for synthesizing composite populations.

2. NGSL Calibration Challenges

Despite STIS being a relatively mature HST instrument when the project began, several calibration issues immediately arose because of the demands and observing mode of this project.

2.1. Wavelength Scale

The NGSL snapshots were carried out in a somewhat nonstandard fashion. Permission was granted to dispense with automatic wavecals, saving $10 - 15$ minutes per snapshot. All of the targets have known radial velocities, so the wavelength scale can be determined *a posteriori*, assuming that grating positioning uncertainties produce only zeropoint shifts. The parameter SHIFTA1 is determined for each grating for each target by cross correlation with a zero velocity template spectrum, then updated so that the X1D extraction produces a spectrum with the published velocity of the target. The X1D spectrum is then deredshifted to the restframe. Some of the metal poor stars have velocities in excess of 300 km s^{-1}, so this extra step is important to obtain the best wavelength scale (Figure 4).

2.2. E1 Aperture

The E1 aperture location was introduced in Cycle 10. Placed near the edge of the CCD, it was implemented to minimize charge transfer losses, but added to the flux calibration challenges in at least two ways. Because of flat fielding and vignetting/throughput differences relative to the standard, central STIS aperture, corrections to the flux calibration sensitivities were necessary. Charles Proffitt reports elsewhere in this volume on these efforts (Proffitt 2006), now nearly finalized. The other challenge arises from an error in the initial location of the 52X0.2E1 aperture, which was off by about 0.5 pixels, resulting in additional slit losses. To save precious time, the NGSL observations (per the STIS Handbook advice) did not use peak-ups to center targets, so for most of Cycle 10, stars were observed slightly off center in the slit. It is not clear yet just how much this affects the spectral energy distribution shape (which is critical for spectral modeling), but overall throughput is affected at the $5-15\%$ level. An example of the typical flux calibration match/mismatch is shown in Figure 5. Even when two gratings do not agree, a simple scaling appears to be enough to correct the discrepancy, but we are comparing the STIS spectra to stellar models to test this.

2.3. Fringing

The STIS CCD is subject to severe fringing at wavelengths redward of 7000Å. Contemporaneous fringe flats were obtained for every target to allow removal of the fringing. Slightly modified versions of the standard IRAF/STSDAS STIS routines are used to divide out the fringes. The results are improved significantly by subtracting a scattered light component; this component is estimated by fitting a smooth curve to the occluded portions of the detector in the fringe flat images. In most cases, the fringes are removed down to the 1-2% level (Figures 2, 4). The flux calibration and useful data overlap between G430L and G750L would be improved by using the flats to take out the high order response/sensitivity variations at the blue end of G750L; presently, these wiggles are removed from the flats using a

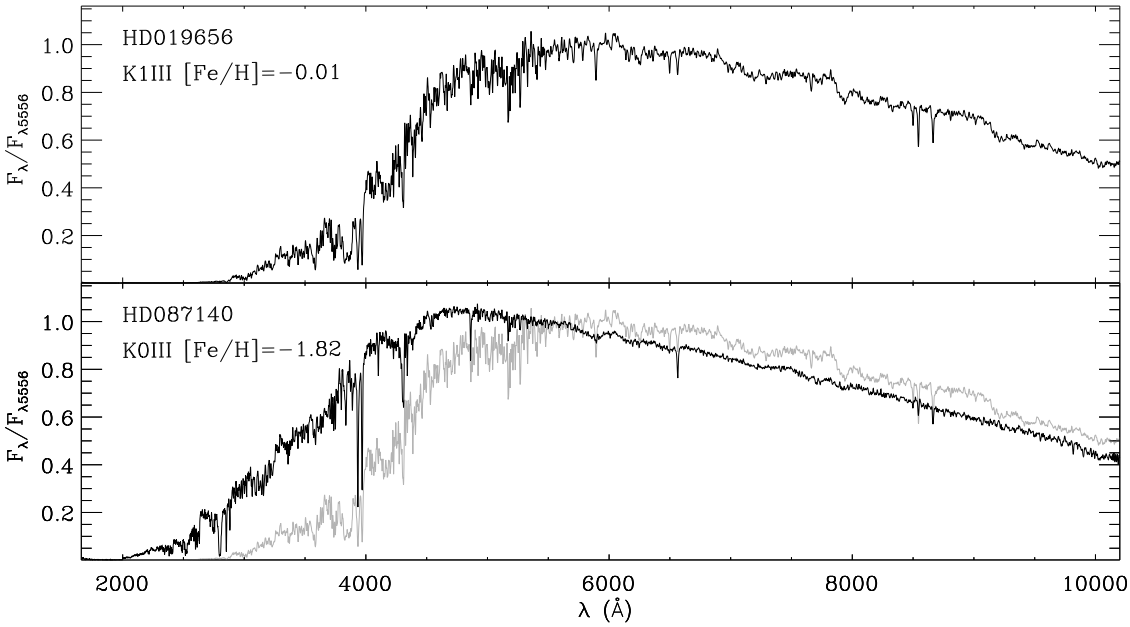

Figure 3: Comparison of STIS spectra of two K giants with very different metallicity. The metal rich star is overplotted as the grey line in the lower panel. This comparison highlights the need for a wide range of not only temperatures but abundances in stellar libraries for proper modeling of composite systems.

very high order spline, necessitating a very high frequency signature in the flux calibration. [1]

2.4. Red Light Contamination in G230LB

The thorniest problem facing NGSL is the red light contamination in the G230LB mode. Because the CCD is extremely sensitive in the red and relatively insensitive in the ultraviolet, even a miniscule fraction of scattered light from red targets can dominate their UV signal. Using the MAMA detectors, which do not suffer from this problem because they are insensitive to optical wavelengths, was not an option because of the brightness of NGSL targets.

The level of the contamination rises from the short to the long wavelength end of G230LB (Figure 6), showing that the main contribution to the contamination is probably dispersed light from G230LB, though some contribution from the wings of the zero-order

[1]Editors' note: Reference files (keyword LFLTFILE) to calibrate large-scale variations in the flats for the CCD L gratings were delivered for use in the pipeline and archive retrieval on 12 December 2005. They provide corrections of up to a few percent at locations beyond the central rows of the detector.

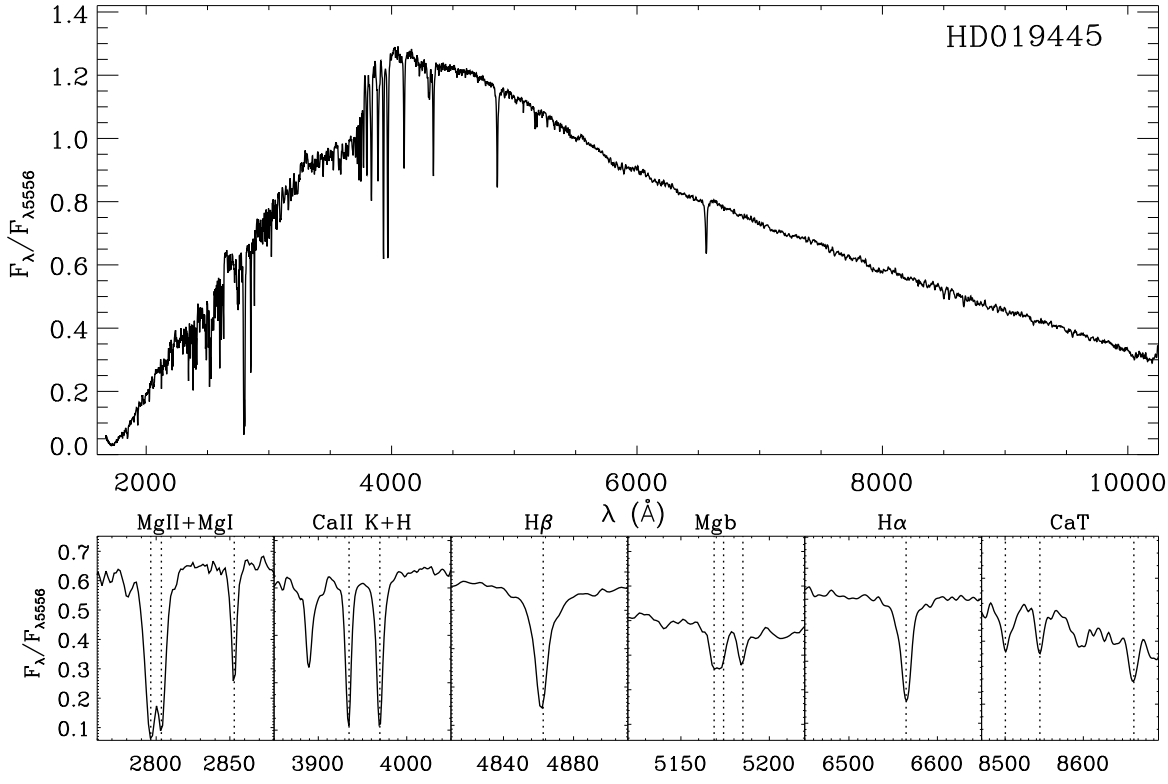

Figure 4: An example NGSL spectrum with all three gratings combined and shifted to restframe wavelengths. The reliability of the wavelength calibration is demonstrated by the relatively good agreement between the absorption features and their expected locations (dotted lines) in the lower panels. The lack of fringing redward of 7000Å demonstrates how well the defringing method works.

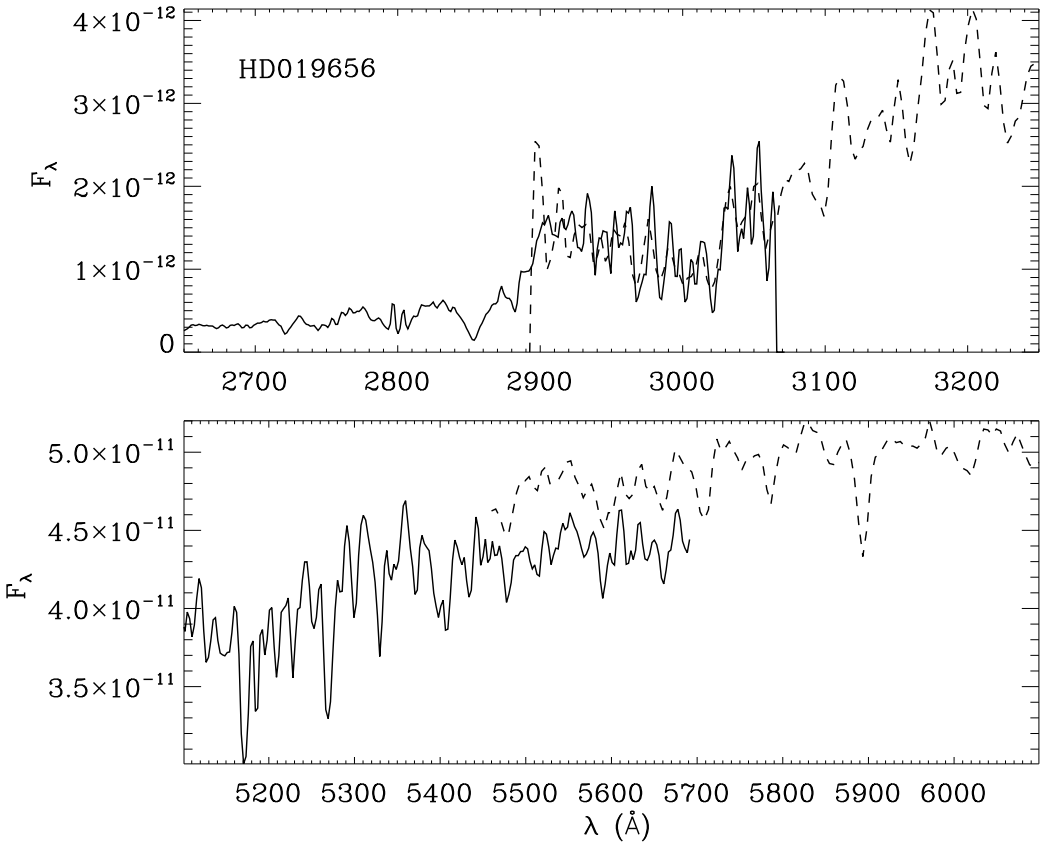

Figure 5: There are sometimes significant mismatches between gratings for the fluxed spectra. For this star, the G430L and G750L spectra differ by 10%, but the G230LB and G430L agree to within a few percent. The disagreement in the red appears to be mainly a scaling or throughput problem and is easily removed, but we have not tested this systematically, nor does this problem occur for all targets. It may be associated with the E1 aperture placement error.

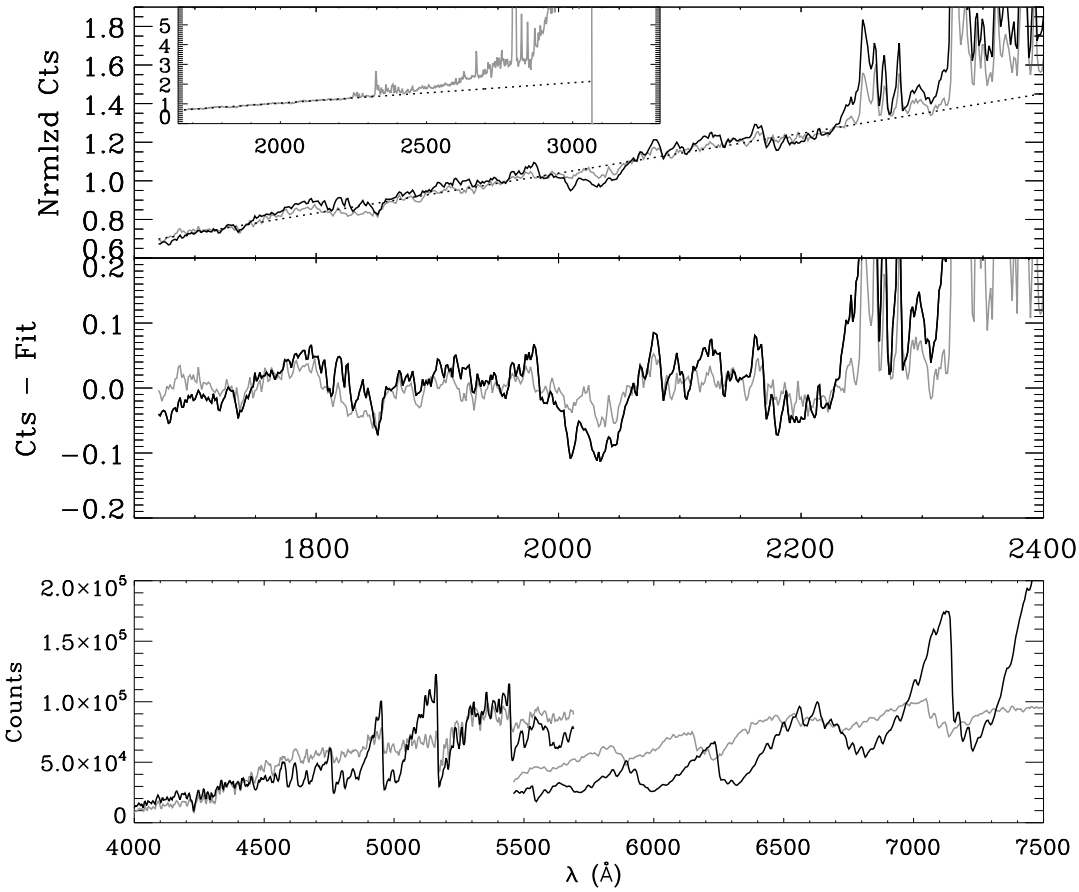

Figure 6: Top panel: Example counts vs. wavelength for two stars to show the rising red light contamination problem for red targets; inset shows the complete G230LB wavelength range. Neither of these stars should generate appreciable counts below 2000Å; all of the detected signal is spurious. Because it rises to the red, the absolute level of contamination is significant at all wavelengths in G230LB. For these stars, normalized to unity at 2000Å, the contamination is remarkably similar in level and slope, allowing removal by a simple linear fit (middle panel). Unfortunately, this breaks down for bluer stars where the correction is still significant but not directly measurable because of actual signal. The bottom panel shows the redder two grating counts for the two stars.

image may also be present. A simple linear fit and subtraction of the contamination works well over a small range of colors (Figure 7), but the slope of the contamination changes with the temperature of the target and there appear to be periodic features in the red light, with scales of a few hundred angstroms (Figure 6). Though scattered light usually does not contain spectral features, there is at least a hint that the later type stars produce scattered light with deeper residuals, perhaps somehow related to the very wide molecular features in the red (Figure 6).

A correction from first principles, or at least a solid empirical calibration for all stellar temperatures, is needed to be able to correct the data adequately. We are at work trying to produce adequate models, cross checking with IUE fluxes.

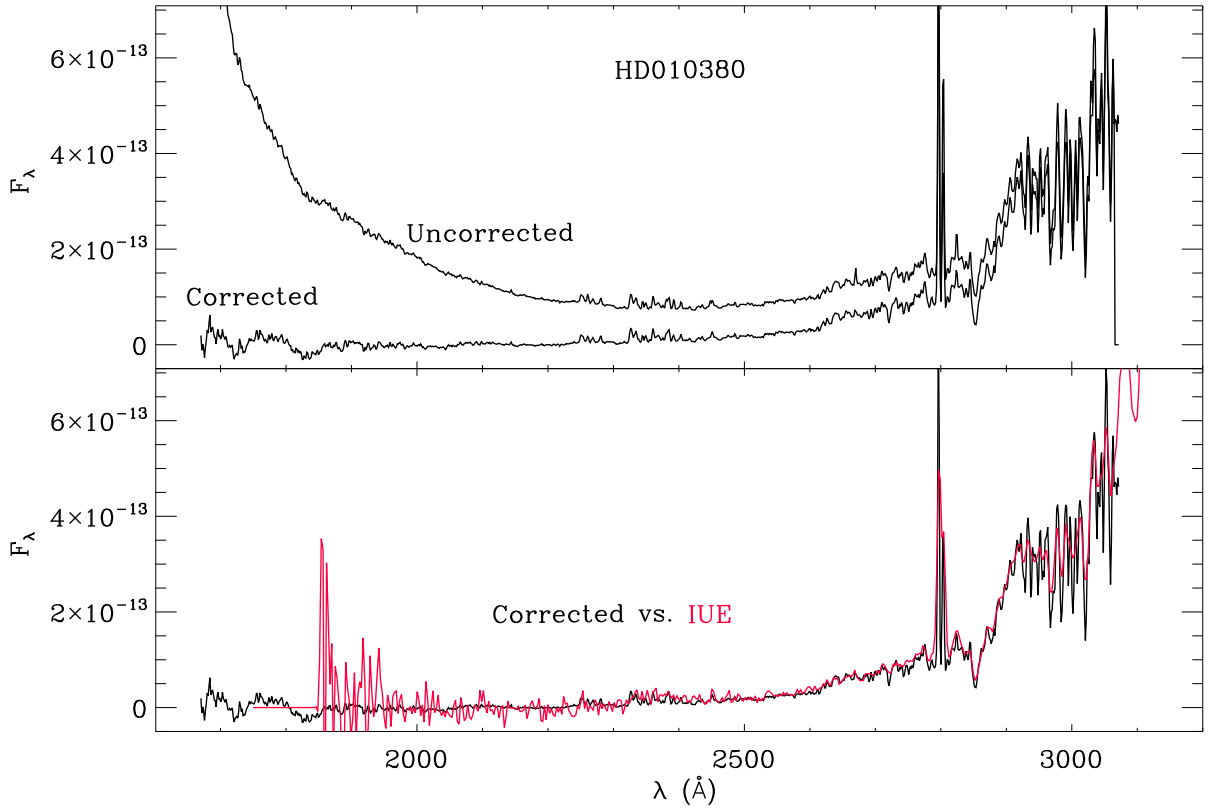

Figure 7: Comparison of IUE and STIS spectra of HD010380 before and after correction for the red contaminating flux. The contamination dominates below 2600Å, and is important at the 5-10% level even at 3000Å.

3. The Future of NGSL

Because the G230LB red light contamination will require more time to model and remove, we are now planning to release the G430L+G750L data on its own by the end of 2005. The UV spectra will be added as soon as a reliable model for the scattered light can be generated. In addition, we plan the following enhancements:

1. Fill out red giant and dwarf color-magnitude space using the HST archive.

2. Extend the wavelength coverage to the far-UV using the HST and IUE archives.

3. We are obtaining echelle resolution spectra using the VLT/UVES instrument. These data will be used to check the atmosphere parameters of the NGSL stars in a consistent manner, plus they will make it possible to work at extremely high resolution over the $3600 - 9000$Å window.

4. Similarly, we have begun extending the spectra to $2.5\mu m$ using the SpeX instrument at IRTF.

5. The final goal of NGSL is to be able to interpolate among the observed stars for spectra of arbitrary $Z - L - T_{eff}$. With 380 instead of 600 input stars, such interpolation becomes somewhat less robust, but with contributions from the archives, it may be possible to reach this goal. And there is always the optimistic view that someday STIS will be repaired and reactivated, and NGSL completed.

The NGSL surpasses all extant compilations in terms of metallicity and wavelength coverage. The spectra, once properly calibrated, will be made publicly available via our website, http://lifshitz.ucdavis.edu/~mgregg/gregg/ngsl/ngsl.html, and other outlets.

Acknowledgments. We gratefully acknowledge generous support from NASA through grants from STScI (GO9088, 9786, and 10222). Much of this work was done at the Institute of Geophysics and Planetary Physics, under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

References

Proffitt, C. R., 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume,199

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The STIS NUV-MAMA Objective Prism and Looking Beyond for HST UV Slitless Spectroscopy

J. Maíz Apellániz[1]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21210, U.S.A.

Abstract. I discuss the recently implemented absolute flux calibration for the STIS objective prism, which yields fluxes accurate to 1% rms in the 1300-3000 Å range relative to the fluxes measured using the first-order modes G140L and G230L. The accuracy of the flux calibration for $\lambda > 3000$ Å is lower. The new calibration includes separate sensitivity curves for the 1200 and 2125 Å settings as well as a time-dependent sensitivity (TDS) solution similar to (but independent of) the one derived for G230L. A new software, MULTISPEC, was used for the data extraction and a brief description is presented. Finally, I also discuss how relevant this analysis is for the future of HST UV slitless spectroscopy.

1. Introduction

This contribution summarizes the work on the STIS objective prism that was presented in two recent ISRs (Maíz Apellániz & Bohlin 2005; Maíz Apellániz & Bohlin 2005).

The objective prism on STIS is a dispersive element with a rather unique characteristic among HST observing modes: spectra can be acquired in a single exposure all the way from 1150 Å to the optical band. This is shown in Fig. 1, where the transformation from wavelength to pixel, as well as the dispersion relation (the inverse of the derivative of the previous function) for the STIS objective prism are shown. The objective prism can use the same slits as the STIS gratings. However, its full power is best revealed in slitless mode, where the compression of the full spectrum of a single source into a small fraction of the horizontal size of the detector can be used to simultaneously obtain spectra of several hundred sources.

Despite those characteristics, the objective prism was not used as frequently in the 7+ years of the STIS lifetime as other spectroscopic observing modes (the archive contains only 123 scientific exposures). The likely explanation is multiple: low sensitivity in the FUV compared to the G140L mode, low spectral resolution in the NUV compared to the G230L or G230LB modes, highly variable spectral resolution, and the difficulties inherent to dealing with crowding in slitless spectral exposures. The first two characteristics relegated the objective prism to some specific scientific tasks; the last two complications were overcome with the availability of specific software for the analysis of STIS objective-prism data.

Another problem is that the existing absolute flux calibration for the prism was known to be worse than that of the gratings due to the small number of spectra on which it was based. Particularly, the repeatability was rather poor. For that reason, the STIS group decided to embark on a project to produce an improved calibration and the results are presented here.

[1] Affiliated with the Space Telescope Division of the European Space Agency, ESTEC, Noordwijk, Netherlands.

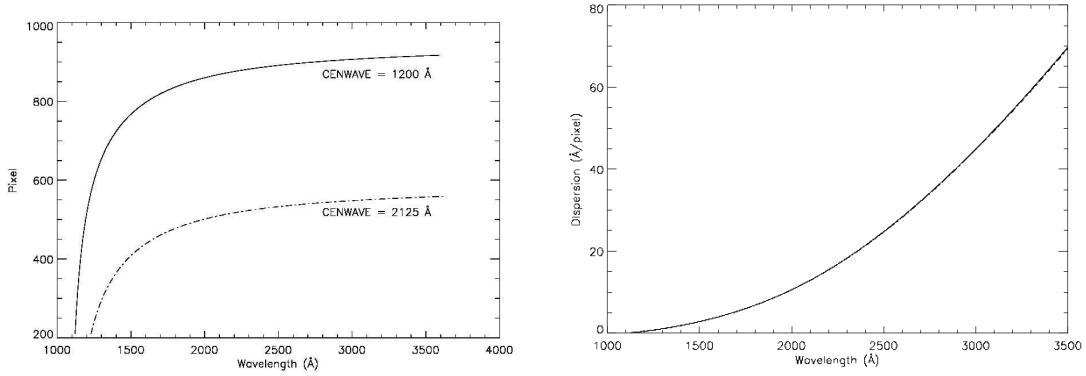

Figure 1: (left) Transformation from wavelength to pixel for the two central wavelength settings in the STIS objective prism. (right) Associated dispersion relation.

2. Data and method

The datasets used for the new flux calibration of the STIS objective prism are observations of the white dwarf standard HS 2027+0651 that cover a large fraction of the STIS lifetime (thus allowing an exploration of time-dependent effects) and were executed using the two available central wavelength settings, 1200 (5 datasets) and 2125 (3 datasets). The wide 52x2 slit was used on 7 datasets and the SRF_2 filter for the eighth. The latter cuts the spectrum below 1280 Å; however, since our analysis will concentrate on $\lambda > 1300$ Å we have also included it. Two of the exposures are shown in Fig. 2.

MULTISPEC (see section 4) was used to extract the spectra by profile fitting in the cross-dispersion direction, one column at a time. The fit residuals were found to contain 1-2% (2125 setting) or 3-5% (1200 setting) of the total flux as a result of the mismatch between the fitted profile and the real one. No significant differences were found between the residuals integrated over an 11-pixel vertical box and those integrated over a 61-pixel one. The 11-pixel residual was added to the spectra obtained by profile fitting to improve the precision of our measurements. A manual shift in the zero x position of the spectrum was applied to each dataset by comparing the FUV absorption lines with those of a reference spectrum of the same star obtained from G140L+G230L data. The obtained displacement was of the order of 1 pixel in all cases. The reason why such a shift is necessary is related to the wavelength-dependent dispersion of objective-prisms. This is illustrated in Fig. 3: if the exact zero x position is not known within better than 1 pixel, then errors of a few percent in the measured fluxes are expected for 1300 Å $< \lambda <$ 3000 Å. The situation is even worse at longer wavelengths, an issue that we will explore in the next section.

After the spectra were extracted, we compared the measured fluxes with those of the reference spectrum.

3. Results

3.1. 1300 Å $< \lambda <$ 3000 Å

We start by analyzing the wavelength range 1300 Å $< \lambda <$ 3000 Å. We plot in Fig. 4 our fluxes for each of the eight datasets divided by the reference spectrum of HS 2027+6051, i.e. the correction that needs to be applied to the sensitivity curve. Results have been rebinned onto a uniform wavelength grid and smoothed to increase the S/N ratio. There are considerable departures from unity as well as differences between different datasets. On a closer look, however, several patterns can be seen:

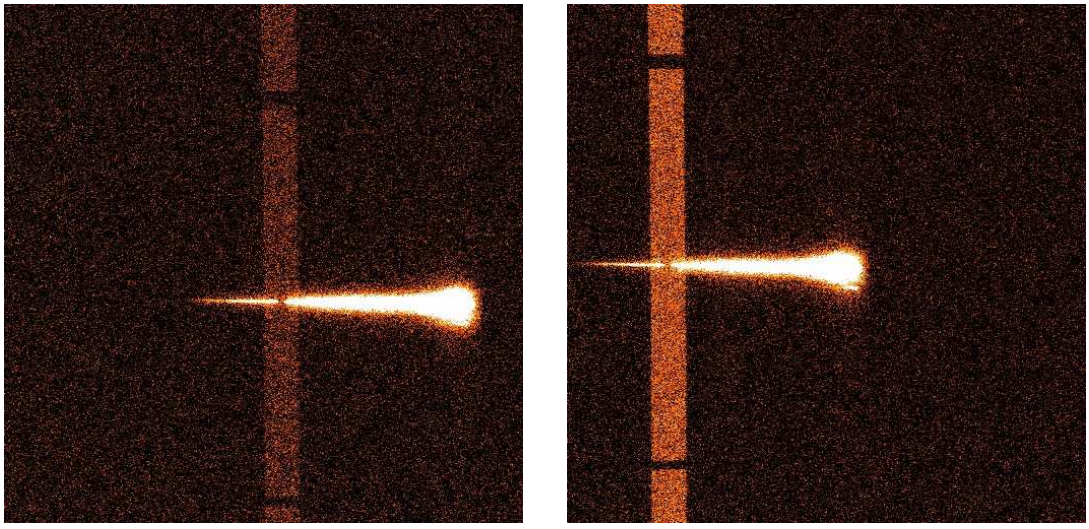

Figure 2: Two of the STIS objective-prism exposures of the standard star HS 2027+6051 used in this work. The left frame shows dataset o8v540060 (1200 setting) and the right frame shows dataset o8v540050 (2125 setting). The vertical band in both cases corresponds to the geocoronal Lyman α emission projected on the 52x2 slit.

- The correction for each of the two settings can be approximately characterized by a single function of wavelength displaced by a constant which is different for each exposure.
- The corrections for the two settings also have similar relative shapes (i.e. the two functions of wavelength have similar characteristics), with a maximum around 2100 Å and a minimum around 2600 Å. However, the two curves appear not to be entirely identical.
- The displacement between each individual exposure in a setting is strongly correlated with the observation date. The sense and magnitude of the correlation with epoch appear to be similar for the two settings.

Those properties led us to devise a correction for the measured fluxes of the form $f_s(\lambda) \times g(t)$, that is, the product of a wavelength-dependent function f (different for each setting s) and a time-dependent function g. Here we will call the first part the sensitivity correction and the second one the time-dependent sensitivity (or TDS) correction. $f_s(\lambda)$ was derived by dividing each of the curves in Fig. 4 by its wavelength-averaged mean and then averaging the resulting functions. $g(t)$ was derived by interpolating in time between the wavelength-averaged means of each of the curves in Fig. 4. Results are shown in Fig. 5.

As expected from our previous impressions, the two sensitivity corrections are similar but not identical. Regarding the TDS correction, the evolution seems to be close to linear in time, with a total variation of $\approx 10\%$ between early 1998 and late 2003. We also tried to detect a possible wavelength dependence on $g(t)$, an effect which is seen for G230L data (Stys et al. 2004), by dividing our data in two ranges, "FUV" (1350-1950 Å) and "NUV" (2000-3000 Å). As seen in Fig. 5, the wavelength dependence is very weak, if present at all. Therefore, we adopted a change of sensitivity with time independent of wavelength[1].

[1]Due to implementation issues, the TDS correction actually used within the STIS pipeline correction differs slightly from the one discussed here in that the eight points in the right panel of Fig. 5 were fitted with a

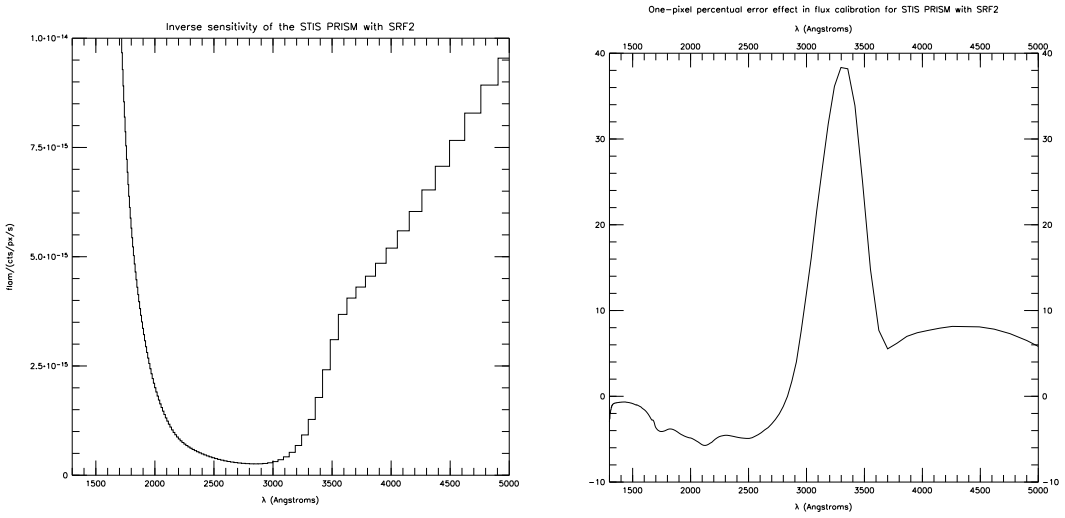

Figure 3: (left) Inverse sensitivity of the STIS objective prism with the SRF_2 filter (changes with respect to the CLEAR mode are minimal in the plotted range). A histogram is used to illustrate the variable nature of the dispersion, with each step corresponding to one pixel. (right) Percent flux error generated by shifting the location of the spectrum by one pixel in the x direction.

The result of applying the new sensitivity calibration is seen in Fig. 6. The rms accuracies are approximately 1% (slightly higher for the 1200 setting, slightly lower for the 2125 setting) and 90% of the data points show accuracies of 2% or better.

The TDS correction developed here has the same sign as the G230L one, and its magnitude is also comparable. At 2000 Å, the G230L TDS is quite similar to the prism one, while at 2800 Å it is $\approx$ 50% that of the prism. We do not have enough temporal coverage to detect the rise and decline of the sensitivity that was found by Stys et al. (2004) in G230L during the first three years of operations of STIS.

3.2. $\lambda >$ 3000 Å

The STIS objective prism can detect photons in the optical band, though at a reduced sensitivity and spectral resolution (see Fig. 3 and note that it is possible to detect photons even with $\lambda >$ 5000 Å). However, two new problems are present for $\lambda >$ 3000 Å that hamper a proper calibration of the data.

The first problem is most severe in the 3000-3500 Å range and has already been mentioned: a mismatch in the zero x position of the spectrum of 1 pixel introduces an error of a few percent at shorter wavelengths; however, it can have disastrous consequences in the 3000-3500 Å range because the slope of the sensitivity curve (as a function of pixel position, not of wavelength) is very large. As seen in the right panel of Fig. 3, the errors induced by such a small displacement can be as high as several tens of percentage points.

The second problem is severe for most of $\lambda >$ 3000 Å and is indirectly related to the previous one. The large slope of the sensitivity curve as a function of pixel position yields large differences in the number of counts between close positions in the detector i.e. the quantity |counts(pixel a)-counts(pixel b)|/distance(a,b) becomes large for small values of distance(a,b). Therefore, one expects that a significant fraction of the counts detected

single straight line instead of using piecewise interpolation. This change raised the formal RMS uncertainty only very slightly (by <0.1%).

around e.g. 3500 Å do not originate from real photons at those wavelengths but from the wings of the line-spread function (LSF, i.e. the component of the PSF in the wavelength direction) at shorter-wavelengths (e.g. from a wavelength of 3000 Å, just a few pixels away, where the number of counts is so much higher that the wings of the LSF, parts of which fall in the pixels of the detector that correspond to 3500 Å, contain a number of counts comparable to the core of the LSF at that longer wavelength). Indeed, it can be seen that the observed cross-dispersion profile broadens considerably around 3500 Å, as expected from such LSF effects.

There is no way to precisely correct these problems, especially the latter because it arises from the intrinsic nature of the LSF. However, we devised an approximate method which consists of:

- As previously stated, using FUV lines such as Lyman α (where the spectral resolution is much better than in the NUV) to fix the zero x position of the spectrum to better than 1 pixel.
- Using the cross-dispersion profile for 3000 Å for all longer wavelengths in the MULTISPEC fit. This is done in order to minimize the artificial broadening caused by the LSF effects.
- Applying the sensitivity and TDS corrections derived in the previous subsection for $\lambda = 3000$ Å to the full optical range redward of 3000 Å.
- Adding the fit residual to the flux in the same way we did in the UV. Several values for the size of the vertical box were tested and the current default of 11 pixels turned out to give the best result.

We plot in Fig. 7 the results obtained for the eight datasets used in this ISR. In the 3000-6000 Å range, the ratio of the prism to the reference flux has a mean of 0.97 and an rms scatter of 0.13, which reflects the precision of our flux calibration in this wavelength range.

4. MULTISPEC: Crowded-field slitless spectroscopy

MULTISPEC is a profile-fitting crowded-field slitless-spectroscopy code written in IDL. It was originally designed specifically for the STIS NUV objective prism but I am currently working on a new version that can be applied to other prisms or gratings in HST instruments. The code is available from the web page `http://www.stsci.edu/~jmaiz` and is described in detail in Maíz Apellániz (2005).

The goal of extracting multiple spectra from a slitless exposure is, in several ways, similar to that of extracting the photometry of a collection of point sources in an image. In the latter case, when one is analyzing a sparse field the most straightforward method is aperture photometry with a small radius. However, for a crowded field the aperture photometry of a star can easily be contaminated by counts produced by neighbors, and a profile-fitting (sometimes called crowded-field) code such as DAOphot, Dophot, or HSTphot does a better job. For slitless spectroscopy, crowding is a larger problem than for photometry because a given source produces a significant number of counts over a larger number of pixels. Hence, given the same number of sources, chances of having an overlap are higher. Therefore, I decided to use profile fitting as the technique for MULTISPEC instead of aperture extraction, the spectroscopic equivalent of aperture photometry.

There are several issues which need to be addressed in fitting profiles to multiple spectra:

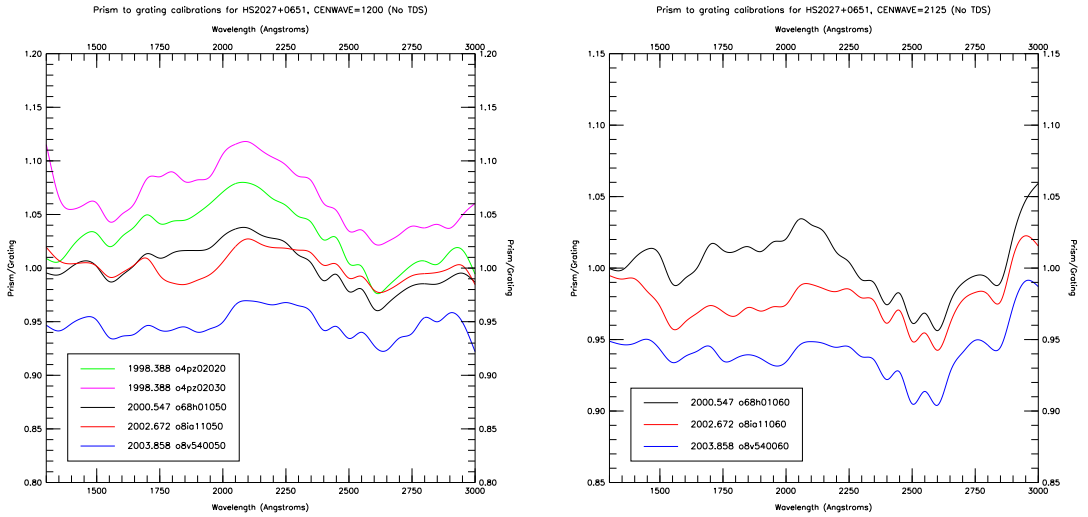

Figure 4: Ratio of the smoothed measured flux divided by the reference (grating) value for the eight datasets in our sample using the existing flux calibration. The left panel shows the five 1200 datasets and the right panel the three 2125 ones.

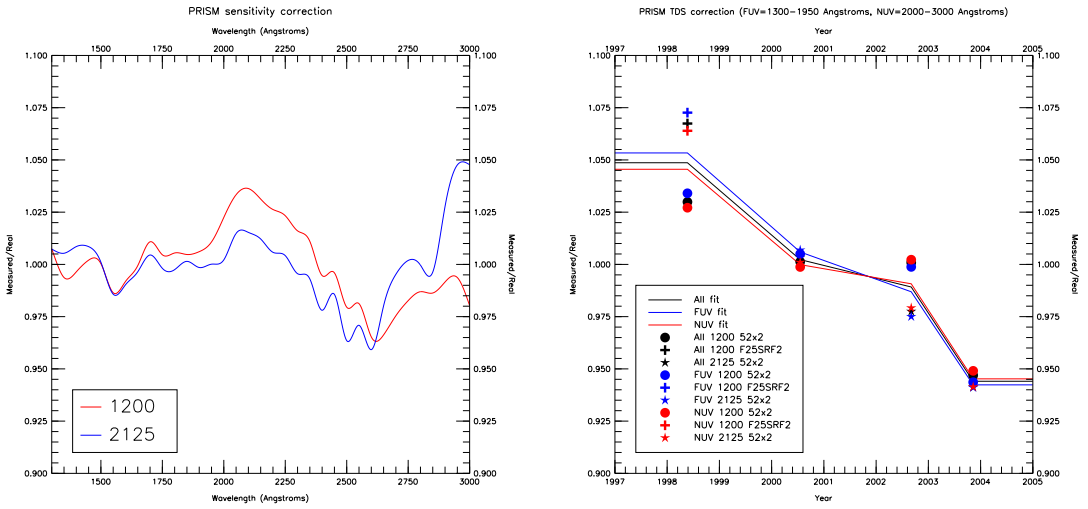

Figure 5: Sensitivity correction (left) and TDS correction (right) for the STIS objective prism derived in this ISR. The left panel shows the two independent solutions derived for each central wavelength setting. The right panel shows the solutions derived for (a) the whole wavelength range [black], (b) 1300-1950 Å [blue], and (c) 2000-3000 Å [red].

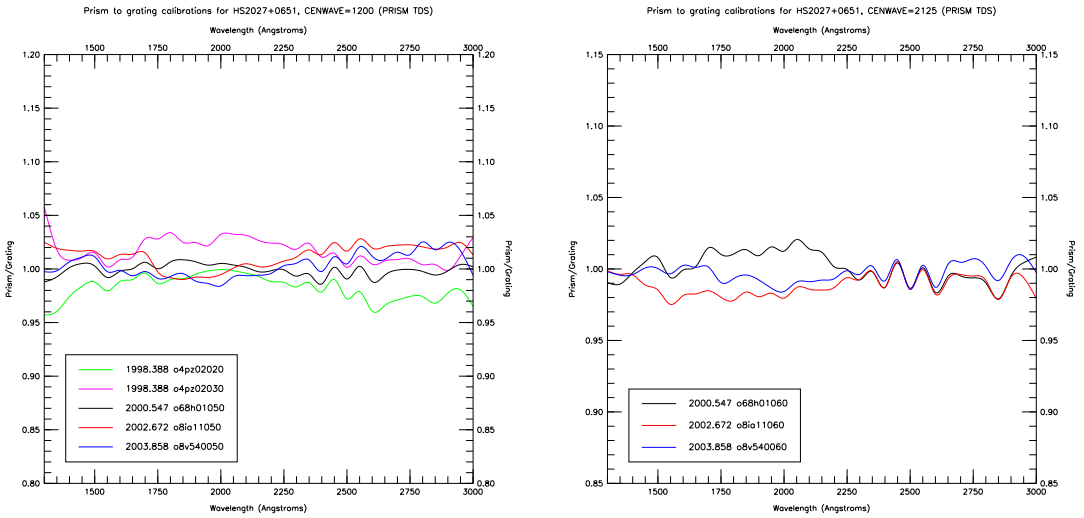

Figure 6: Same as Fig. 4 after applying the sensitivity correction and the TDS solution derived in this work.

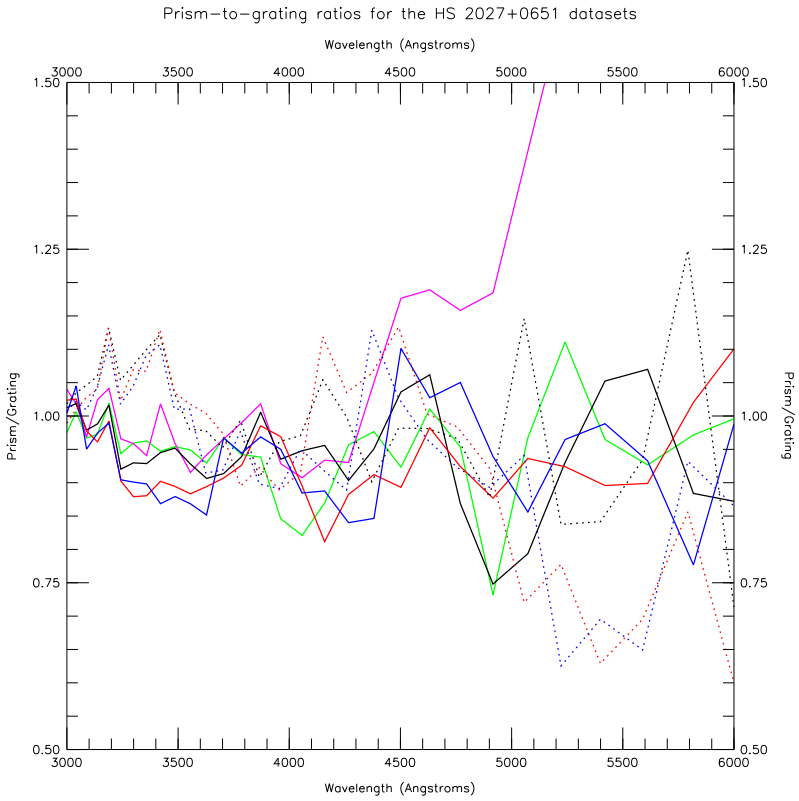

Figure 7: Prism to reference flux ratios in the 3000-6000 Å range for the eight HS 2027+0651 datasets shown in Fig. 6. The color coding is the same one as in that figure, with solid lines for the 1200 setting and dotted ones for the 2125 setting.

- A difference between a photometry- and a spectroscopy-profile-fitting code is that the former has no preferred direction in the detector and, for that reason, it uses 2-D profiles as its fundamental unit. Spectroscopic exposures, on the other hand, have a distinct asymmetry in that one direction shares the dual character of being an angular – or spatial – and a wavelength – or spectral – coordinate (that direction will be called here the dispersion direction). Furthermore, the position of the object in the perpendicular angular coordinate (the cross-dispersion direction) is fixed over its extent. For those reasons, it makes more sense to fit 1-D profiles in the cross-dispersion direction once we have determined the position of the object on the 2-D detector (i.e. its trace).

- Another difference between the two types of code is the need for auxiliary exposures. A photometry package can use one image to identify the sources and to extract the magnitudes of the objects. Doing so in a spectroscopy code is more difficult, given the possible extensive overlap among sources. For that reason, an image (or images) of the same field is a preferred alternative for source identification.

- The use of two different types of observations, an image (or images) for source identification and a slitless spectral exposure for the measurement of spectra, introduces an added complication: the need to accurately know the relative geometric distortion between the two in order to produce a one-to-one correspondence between the positions in one and the other throughout the field covered. These geometric distortions can be calculated by obtaining calibration spectral slitless exposures and images of a crowded field. The images are used to generate a simulated slitless exposure of the field without any distortion which can then be compared with the real exposure. Note that, due to the degeneracy between spatial and spectral coordinates in the dispersion direction, the precision of the geometric distortion solution in that direction is likely to be poorer than in the cross-dispersion direction unless the presence of absorption/and or emission lines is ubiquitous in the sources, their intensities and wavelengths are well known, and the dispersing element can resolve them.

- As is also the case for crowded-field photometry, a detailed knowledge of the spatial (cross-dispersion) profile is required for a crowded-field spectroscopy code. Such a profile is likely to be wavelength dependent. If the spatial profile is not well known, systematic errors can be introduced, which would be small for bright stars and caused by an incorrect aperture correction but possibly large for dim stars, especially those close to bright ones.

- The ability of a profile-fitting code to accurately deconvolve the fluxes of a close pair is expected to depend heavily on the S/N ratio of the data (Porter et al. 2004). Therefore, we should expect different uncertainties for stars which have the same number of counts at the same wavelength depending on the presence or absence of neighbors.

MULTISPEC has been written with all of the above considerations is mind. It is a modular package, allowing the user to execute the different steps such as background calculation or profile fitting with different parameters and to iterate over them in order to improve the quality of the results. The reader is referred to Maíz Apellániz (2005) for further details but in Figure 8 I show an example of MULTISPEC inputs and outputs.

5. Conclusions and the future of UV slitless spectrocopy with HST

We have developed new photometric calibrations for the two central wavelength settings of the STIS objective prism mode. We have also calculated a wavelength-independent

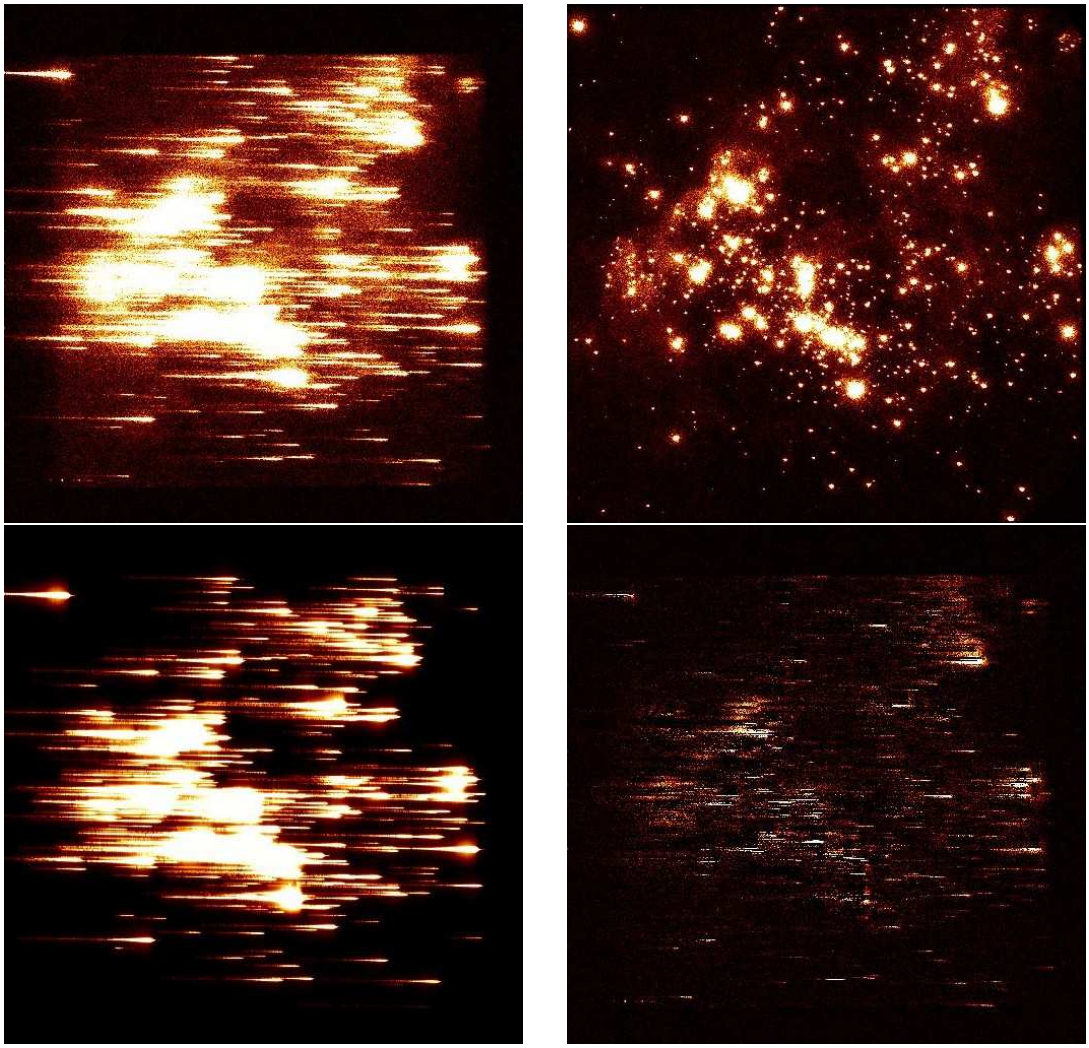

Figure 8: Different 2-D data used or generated by MULTISPEC for a crowded field in NGC 604. [upper left] Spectral exposure. [upper right] Associated exposure in imaging mode. [lower left] Modeled data. [lower right] Full residual (exposure minus modeled data and background). All panels use the same color scale and count limits. Note that for the STIS NUV-MAMA the plate scale (in arcseconds/pixel) for an image is 83.4% that of an objective-prism spectral exposure (compare the data edges in the two upperpanels).

TDS correction specifically for the prism modes. The combination of both corrections allows for an rms accuracy of 1% with respect to the G140L+G230L results in the 1300-3000 Å range, which themselves have an absolute accuracy of 3-4% (Bohlin 2000). For $\lambda > 3000$ Å, the accuracy is lower due to imprecisions in the assignment of wavelengths to the spectrum on the detector and to contamination between adjacent wavelengths.

If STIS is revived later on, the calibration presented here could be used for future observations with the NUV objective prism. A reborn STIS would also provide opportunities for slitless spectroscopy of crowded fields with gratings such as G140L, maybe combined with filters such as F25QTZ or F25SRF2 to reduce the extension of the traces on the detector and thus avoid confusion. In either case, MULTISPEC could be applied for the extraction of spectra from those data.

In the meantime, UV spectroscopy with HST can still be done in the FUV with SBC using either PR110L or PR130L or in the NUV with the HRC using PR200L. The two SBC prisms provide access to the FUV with a better efficiency than the STIS objective prism while PR200L is restricted to the NUV but has the advantage over the STIS objective prism of not being subjected to bright-object limits. The ST-ECF and STScI ACS groups are in the process of obtaining refined calibrations for those prisms through calibration programs 10391, 10722, 10736, and 10743 (see contributions in these proceedings by Martin Kümmel and Søren Larsen). In particular, we are obtaining exposures of crowded fields with all three prisms to generate the calibration files required by MULTISPEC.

References

Bohlin, R. C. 2000, *AJ* **120**, 437

Maíz Apellániz, J. 2005, *Instrument Science Report* STIS 2005-02 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

Maíz Apellániz, J. & Bohlin, R. C. 2005, *Instrument Science Report* STIS 2005-01 (Baltimore: STScI)

Porter, J. M., Oudmaijer, R. D., & Baines, D. 2004,*A&A* **428**, 327

Stys, D. J., Bohlin, R. C., & Goudfrooij, P. 2004, *Instrument Science Report* STIS 2004-04 (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

A Recalibration of Optical Photometry Based on STIS Spectrophotometry

J. Maíz Apellániz[1]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21210, U.S.A.

Abstract. I use STIS spectrophotometry to test the published sensitivity curves $P(\lambda)$ and calculate the zero points ZP_P for three optical photometry systems: Tycho-2 $B_\mathrm{T}V_\mathrm{T}$, Strömgren $uvby$, and Johnson UBV. The analysis includes the new recalibration of grating/aperture corrections, vignetting, and charge transfer inefficiency effects produced by the STIS group and considers the consequences of both random and systematic uncertainties. For Tycho-2 B_T and V_T, Strömgren v, b, and y, and Johnson B and V, I find that the published sensitivity curves are consistent with the available photometry and spectrophotometry and I provide new values for the associated zero points. The situation is different for the Strömgren u and the Johnson U filters. There I find that the published sensitivity curves yield results that are inconsistent with the available photometry and spectrophotometry, likely caused by an incorrect treatment of atmospheric effects in the short-wavelength end. I reanalyze the data to produce new average sensitivity curves for those two filters and new values for ZP_{c_1} and ZP_{U-B}. The new computation of synthetic $U-B$ and $B-V$ colors uses a single B sensitivity curve, which eliminates the previous unphysical existence of different definitions for each color.

1. Description

- I use STIS spectrophotometry to test the published sensitivity curves $P(\lambda)$ and calculate the zero points ZP_P for three optical photometry systems.

$$m_P = -2.5 \log_{10} \left(\frac{\int P(\lambda) f_\lambda(\lambda) \lambda \, d\lambda}{\int P(\lambda) f_{\lambda,\mathrm{ref}}(\lambda) \lambda \, d\lambda} \right) + \mathrm{ZP}_P$$

- The spectrophotometry was obtained from (a) the Next Generation Spectral Library (NGSL, Gregg et al. 2004 and this meeting) which includes 255 non-variable good-quality stars and from (b) the Bohlin sample (Bohlin et al. 2001, Bohlin & Gilliland 2004a) which includes another 19 stars. The two samples span a wide range of temperatures, gravities and metallicities. For each of the 274 stars a high S/N spectrum that includes the 3 000-7 000 Å wavelength range was built by combining data obtained with the STIS G430L and G750L gratings. The data were processed using the latest calibration files and algorithms from the STIS pipeline.

- The Tycho-2 photometry for the NGSL and Bohlin samples was collected directly from Høg et al. (2000) while the equivalent Strömgren and Johnson photometry had

[1]Affiliated wuth the Space Telescope Division of the European Space Agency, ESTEC, Noordwijk, Netherlands.

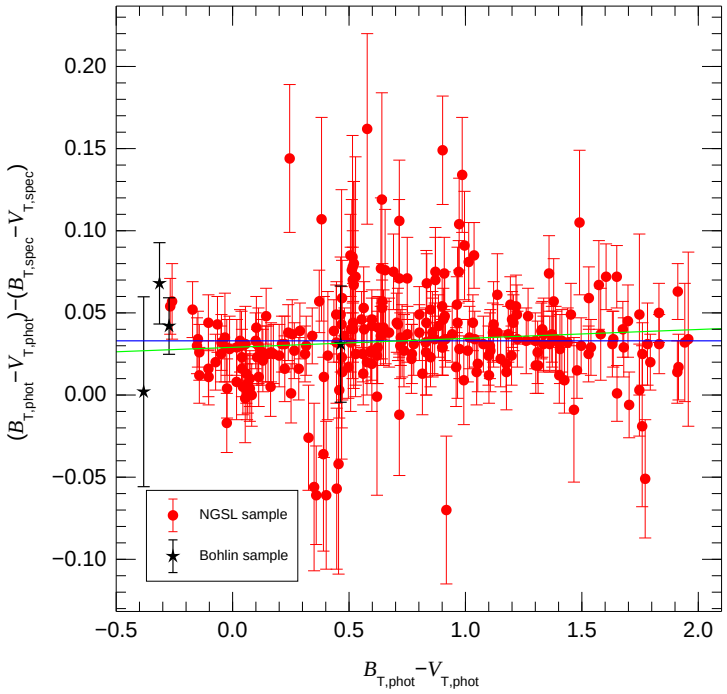

Figure 1: Comparison between photometric and spectrophotometric $B_T - V_T$ colors as a function of photometric $B_T - V_T$ for the NGSL + Bohlin samples using the Bessell (2000) sensitivity curves. The error bars represent the photometric uncertainties and the horizontal blue line marks the proposed $ZP_{B_T-V_T}$. The green line shows the result of a weighted linear fit to the data.

to be compiled from the General Catalogue of Photometric Data (GCPD, Mermilliod et al. 1997). A careful procedure to eliminate incorrect photometric values and to assign more significance to the data with better quality was followed in the latter case.

- The full results of this work are presented in Maíz Apellániz (2006). The ZPs and the new sensitivity curves (see below) will be implemented in the version 2.0 of CHORIZOS (Maíz Apellániz 2004)

2. Tycho-2 B_TV_T

- I have reanalyzed the previous results of Maíz Apellániz (2005), which were obtained using a previous STIS calibration.
- The B_T and V_T sensitivity curves of Bessell (1990) are accurate, as evidenced by the small slope present in the fit in Figure 1. This confirms the previous results on the stability for the Tycho-2 photometry and the adequacy of the NGSL and Bohlin datasets for the testing of UV-optical sensitivity curves and the calculation of ZPs.
- New ZPs are presented in Tables 1 and 2 which are only slightly different from those given by Maíz Apellániz (2005).
- The consistency of the V_T calibration was checked by computing ZP_{V_T} using two different datasets and methods.

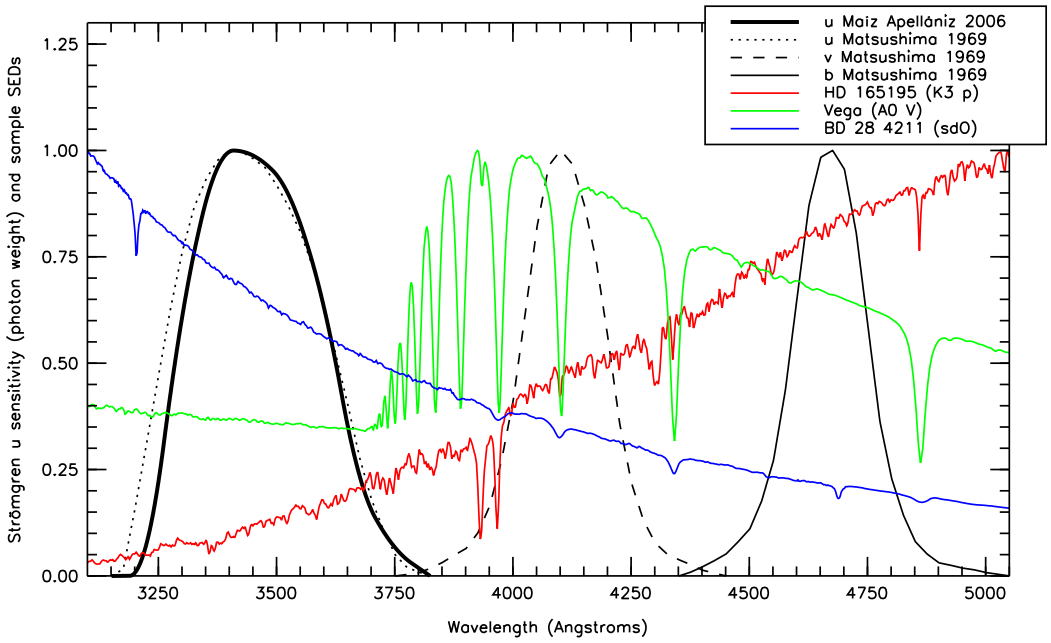

Figure 2: Matsushima (1969) photon-counting normalized sensitivity curves for the the three Strömgren filters u, v, and b filters and the new Maíz Apellániz (2006) curve for u. Three normalized SEDs from the two samples used to calibrate the Strömgren photometry are also shown as f_λ.

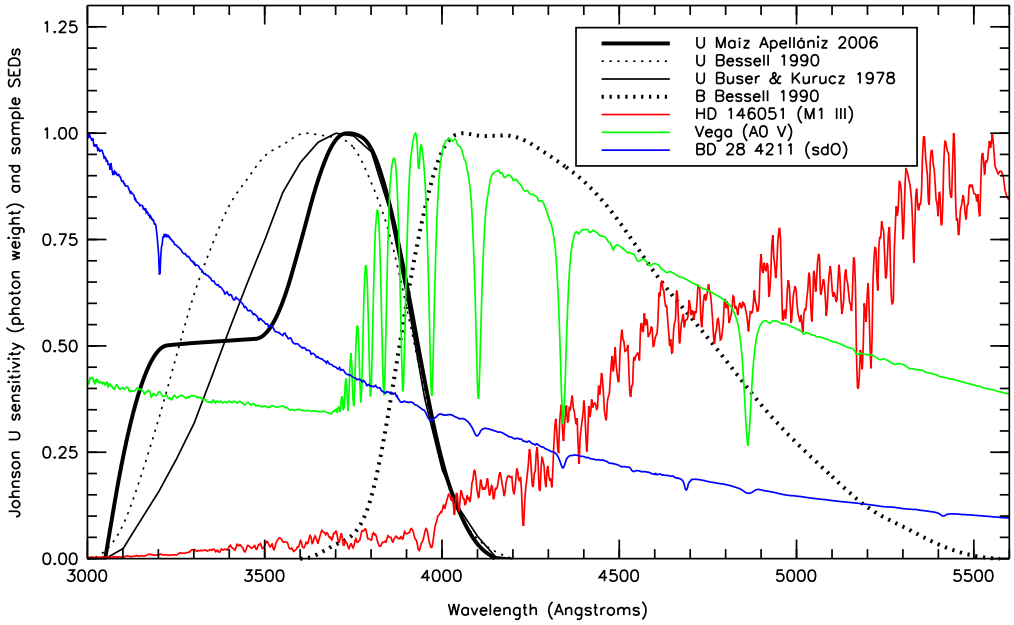

Figure 3: Bessell (1990) U (UX) and B (B) and Buser & Kurucz (1978) U (U_3) photon-counting normalized sensitivity curves and the new Maíz Apellániz (2006) curve for U. Three normalized SEDs from the two samples used to calibrate the Johnson photometry are also shown as f_λ.

- The Tycho-2 ZPs have precisions better than one hundredth of a magnitude, implying that the published uncertainties (which are typically several times larger) can be in most cases directly compared with synthetic photometry results without having to invoke additional sources of systematic errors.

3. Strömgren *uvby*

- I have collected $b-y$, m_1 $(v-2b+y)$, and c_1 $(u-2v+b)$ data from the GCPD. After the selection and filtering process, I ended up with a total of 104 stars, 100 from the NGSL sample and 4 from the Bohlin sample.

- The v, b, and y sensitivity curves of Matsushima (1969) are accurate.

- The u sensitivity curve of Matsushima (1969) yields a significant slope in a $(b-y)_{\rm phot}$ vs. $c_{1,\rm phot}$-$c_{1,\rm spec}$ plot, indicating that it is not an accurate representation of the published photometry.

- I have used a χ^2 minimization algorithm to calculate a new Strömgren u sensitivity curve (Figure 2). The new curve is quite similar to the Matsushima (1969) one but has a short-wavelength edge which is redder by ≈ 35 Å. The likely culprit for the difference between the two is a difference in atmospheric transmission or in the reduction process to extrapolate to zero air mass. The new sensitivity curve yields a much better fit to the data.

- The new ZPs are given in Table 1. Their precisions are similar to those for Tycho-2.

- I show in Table 3 the typical uncertainties for individual Strömgren photometric entries in the GCPD, as measured by comparing them with the corresponding spectrophotometric values. Little variation is observed as a function of magnitude or the number of measurements per entry, so those values can be applied to the full GCPD once the outliers (which comprise 2-6% of the entries) are eliminated. In all three cases ($b-y$, m_1, and c_1) the uncertainties are several times larger than the ZP precisions, so the data can be directly compared with synthetic photometry results.

4. Johnson *UBV*

- I have collected V, $B-V$, and $U-B$ data from the GCPD. After the selection and filtering process, I ended up with a total of 108, 111, and 101 stars, respectively.

- The data are consistent with the Vega V magnitude of 0.026 ± 0.008 derived by Bohlin & Gilliland (2004b), which can be used as $\rm ZP_V$.

- The B and V sensitivity curves of Bessell (1990) are accurate.

- Both of the U sensitivity curves of Buser & Kurucz (1978) and Bessell (1990) are inconsistent with the observed $U-B$ photometry.

- I have derived a new Johnson U sensitivity curve (Figure 3). The new curve has a plateau in the 3200-3500 Å region that is likely caused by an averaging over different effective sensitivities used under different conditions. The fit to the $U-B$ data using the new U sensitivity curve is better than using the previous two, but the residuals are larger than for other colors and indices.

- The new ZPs are given in Table 1. The precision for ZP_{B-V} is similar to the previous ones but the one for ZP_{U-B} is considerably worse. Therefore, it is not possible to eliminate systematic errors beyond 0.01 magnitudes for archival $U-B$ photometry.
- I show in Table 3 the typical uncertainties for individual Johnson photometric entries in the GCPD. As in the Strömgren case, little variation is observed as a function of magnitude and the outliers comprise 2-6% of the entries. The random uncertainties for Johnson photometry are found to be larger than for the Strömgren case.

Table 1: Color/index zero points and associated uncertainties/errors.

	Tycho-2	Strömgren			Johnson	
	$B_\mathrm{T}-V_\mathrm{T}$	$b-y$	m_1	c_1	$B-V$	$U-B$
zero point	0.033	0.007	0.154	1.092	0.010	0.020
random	0.001	0.001	0.001	0.002	0.001	0.006
systematic	0.005	0.003	0.003	0.004	0.004	0.014

Table 2: Magnitude zero points and associated uncertainties.

	Tycho-2	Johnson
	V_T	V
zero point	0.034	0.026
uncertainty	0.006	0.008

Table 3: Random uncertainties for individual entries in the GCPD.

Strömgren			Johnson	
$b-y$	m_1	c_1	$B-V$	$U-B$
0.013	0.014	0.018	0.020	0.028

Acknowledgments. I would like to thank Ralph Bohlin and Jay Holberg for useful suggestions.

References

Bessell, M. S. 1990, *PASP* **102**, 1181
Bessell, M. S. 2000, *PASP* **112**, 961
Bohlin, R. C., Dickinson, M. E., and Calzetti, D. 2001, *AJ* **122**, 2118
Bohlin, R. C. and Gilliland, R. L. 2004a, *AJ* **128**, 3053
Bohlin, R. C. and Gilliland, R. L. 2004b, *AJ* **127**, 3508
Buser, R. and Kurucz, R. L. 1978, *A&A* **70**, 555
Gregg, M. D., et al. 2004, *American Astronomical Society Meeting Abstracts* 205
Høg, E., et al. 2000, *A&A* **355**, L27
Maíz Apellániz, J. 2004, *PASP* **116**, 859
Maíz Apellániz, J. 2005, *PASP* **117**, 615
Maíz Apellániz, J. 2006, *AJ* **131**, 1184
Matsushima, S. 1969, *ApJ* **158**, 1137
Mermilliod, J.-C., Mermilliod, M., and Hauck, B. 1997, *A&AS* **124**, 349

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Calibration and Interpretation of STIS Imaging Mode Fluxes

Charles R. Proffitt[1]

Science Programs, Computer Sciences Corporation, Baltimore, MD 21218

Abstract. While primarily a spectroscopic instrument, STIS also had a number of imaging modes, many of which provided unique capabilities. However, most of the STIS imaging modes also had very wide and non-standard passbands, making calibration and interpretation of STIS imaging data especially challenging. Here we describe the methods used to derive the adopted throughput curves, zero points, and aperture corrections for these modes. The wide, non-standard passbands of STIS imaging modes also present special challenges when comparing with observations in other photometric systems. We discuss the relationship between STIS magnitudes and colors and those of other common photometric systems.

1. STIS Imaging Throughput Time Dependence and Zero Points.

The STIS instrument has three detectors, a CCD sensitive from $\sim$ 2000 to 10300 Å, a Cs_2Te Multi-Anode Microchannel Array (MAMA) detector (1600 to 3100 Å) and a solar-blind CsI MAMA (1150 to 1700 Å). Sensitivities of these detectors have varied over time. For a given detector these changes depend on wavelength, but appear to be independent of the mode. So for imaging observations we adopt the same time dependence vs. wavelength that was determined for the first order spectroscopic modes by Stys et al. (2004). This leads to considerable color dependence of the sensitivity changes with time. In Figure 1 we give examples of the predicted sensitivity changes for stars of different colors at 380 day intervals over the 7+ year operational life of STIS.

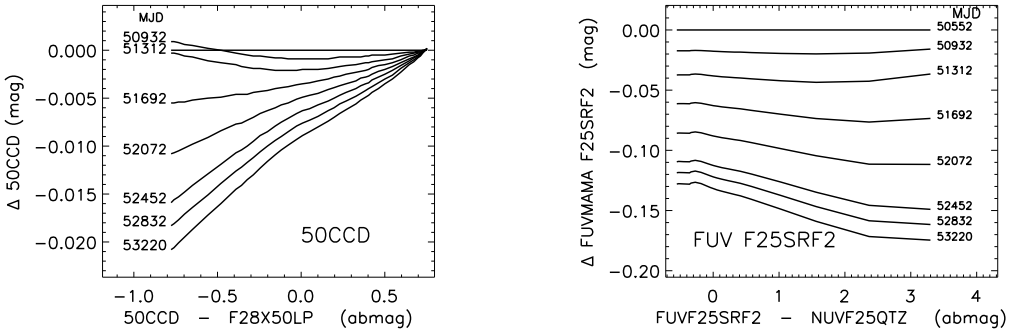

Figure 1: Predicted changes in sensitivity as a function of time and color for selected STIS modes are shown at 380 day intervals over the operational lifetime of the STIS instrument.

The available detector/imaging filter combinations are summarized in Table 1 for a time near the beginning of STIS operations. In this table PHOTPLAM is the pivot wavelength

[1]also Space Telescope Science Institute, and Catholic University of America

in Å, PHOTFLAM is the flux (ergs/cm^2/s/Å) of a flat spectrum that produces 1 cnt/sec in each mode, and PHOTBW the passband RMS width in Å. The corresponding keywords in individual STIS image headers take into account the time dependent sensitivity changes, and contain values appropriate for the actual date of each observation. See the "SYNPHOT User's Guide" (Laidler et al. 2005) for detailed definitions of these parameters.

Table 1: STIS Imaging Mode Zero Points on April 14, 1997

Detector	Aperture	PHOTPLAM	PHOTFLAM	PHOTBW	STmag	ABmag	VEGAmag
CCD	50ccd	5733.23	9.951e-20	1840.03	26.405	26.305	26.061
	f28x50lp	7216.75	1.639e-19	1139.84	25.863	25.264	24.990
	f28x50oii	3738.06	4.690e-17	24.14	19.722	20.551	19.724
	f28x50oiii*	5005.8	1.337e-16	2.65	18.585	18.780	18.864
	f25nd3	6314.84	4.508e-17	1987.74	19.765	19.455	19.197
	f25nd5	7010.84	7.400e-15	2000.83	14.227	13.690	13.401
NUV-MAMA	f25srf2	2304.01	5.567e-18	478.36	22.036	23.916	22.238
	f25qtz	2357.89	5.922e-18	420.17	21.969	23.798	22.161
	f25cniii	2263.35	4.986e-18	508.34	22.156	24.074	22.357
	f25cn182	2006.92	5.932e-17	285.65	19.467	21.646	19.868
	f25mgii	2873.82	2.015e-16	248.98	18.139	19.539	18.151
	f25cn270	2722.84	3.923e-17	121.84	19.916	21.433	19.951
	f25nd3	2358.49	4.461e-15	553.16	14.776	16.605	14.931
	f25nd5	2630.39	3.359e-12	623.34	7.585	9.177	7.660
FUV-MAMA	25mama	1374.38	2.103e-17	137.60	20.593	23.595	20.732
	f25srf2	1457.38	3.800e-17	120.47	19.951	22.825	20.447
	f25qtz	1595.04	1.036e-16	96.74	18.861	21.540	19.516
	f25lya	1243.08	1.048e-15	61.74	16.350	19.569	14.915
	f25nd3	1376.05	2.146e-14	134.27	13.071	16.070	13.229
	f25nd5	1385.09	3.011e-11	141.66	5.203	8.188	5.383

*Calculated parameters for the OIII filter exclude the red leak

2. Determination of Absolute Throughputs

Since STIS is a spectrograph as well as an imaging instrument, the relative throughputs of different filters as a function of wavelength were easily determined by comparing filtered and unfiltered spectra of the same stars. To determine the wavelength dependent throughput and aperture corrections for each detector's unfiltered mode, deep images of a small number of stars with well measured spectral energy distributions (SEDs) and spanning a wide range of colors were observed with several filters. The detector throughputs vs. wavelength were then adjusted to match the observed count rates for each filter and target.

For CCD imaging modes, the details of this calibration, including wavelength dependent aperture corrections, were presented in Proffitt (2004). The overall CCD throughput at long wavelengths had to be increased by up to 18% above prelaunch estimates. In addition, the scattering of light into an extended red halo put up to 20% of the flux at distances from the image center of more than 0.5".

For the broadband MAMA filters, existing standard stars were too bright, and to provide well calibrated standards long-slit STIS G140L, G230L, and G430L spectra were taken of a number of hot HB stars in the globular cluster NGC 6681. This cluster was also used as the primary calibration object for the long term monitoring of the MAMA imaging stability, and so copious sets of images of these same stars are available in most MAMA imaging modes. Details of the MAMA imaging absolute flux calibration are given in Proffitt et al. (2003).

3. Transforming STIS Magnitudes to other Photometric Systems

Given the large differences between STIS bandpasses and commonly used ground-based photometric systems, the relationships between the different systems will be very sensitive to the details of the spectral energy distributions (SED) being considered. The most rigorous way to intercompare between such different bandpasses would be to use a series of spectral energy distributions appropriate for the objects being studied to perform detailed synthetic spectrum calculations for the all photometric bandpasses being used. However, such detailed calculations can be time consuming to set up, and when performing quick evaluations of STIS imaging data it would still be very useful to have some simple transformations between STIS imaging magnitudes and those of more common systems, as long as the limitations of such transformations are understood.

3.1. Relationship of STIS CCD Imaging Bandpasses to other Photometric Systems

Comparing the STIS CCD throughput curves to standard broadband Johnson filters shows that the STIS passbands (Figure 2) extend to significantly longer and shorter wavelengths than does the $UBVR_cI_c$ system, and any transformation between STIS and other broadband photometric systems will depend strongly on the details of the SED. This is illustrated in Figure 3, where we plot synthetic colors and magnitudes calculated for several sequences of Kurucz main-sequence models with different metallicities and assumed reddenings. Each sequence ranges in $T_{\rm eff}$ from 50000 to 3500 K.

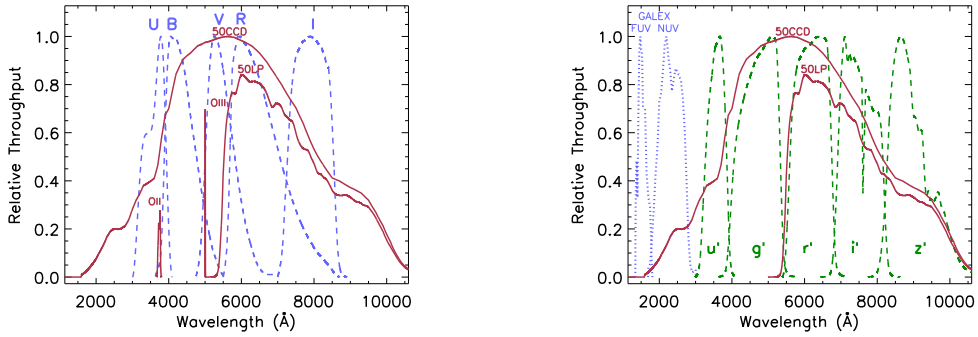

Figure 2: The relative throughput of the STIS CCD imaging modes (solid lines) is compared to that of Johnson bands (dashed lines) in the left hand figure, and in the right hand figure to the throughput of SDSS (dashed lines) and GALEX (dotted lines) bandpasses.

Color-color relations (Figure 4) are somewhat tighter, especially when comparing STIS CCD colors to $B - I_c$, but differential reddening still has a strong effect on the transformation.

The combined long wavelength filters of Gunn type photometric systems, such as that used for the Sloane Digital Sky Survey (York et al., 2000), provide a better match to the F28X50LP filter throughput (right hand panel of Figure 2). While no ground based system can reach wavelengths as short as the unfiltered STIS CCD's cutoff, these missing UV wavelengths are filled in by the GALEX NUV band. This makes it possible to use SDSS plus GALEX photometry to robustly predict STIS CCD magnitudes. This may be useful to place STIS observations in the context of these surveys or to easily use synthetic colors calculated for these systems when interpreting STIS imaging observations.

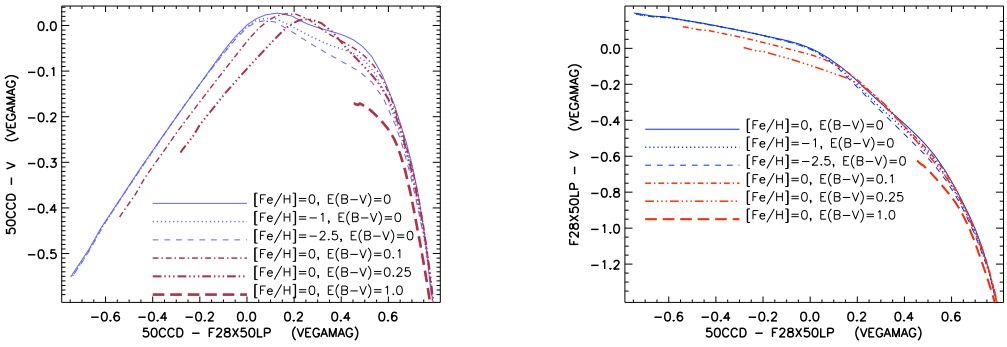

Figure 3: Relation between STIS 50CCD (left) and F28X50LP (right) magnitudes and Johnson V band as a function of color for six series of Kurucz main-sequence models.

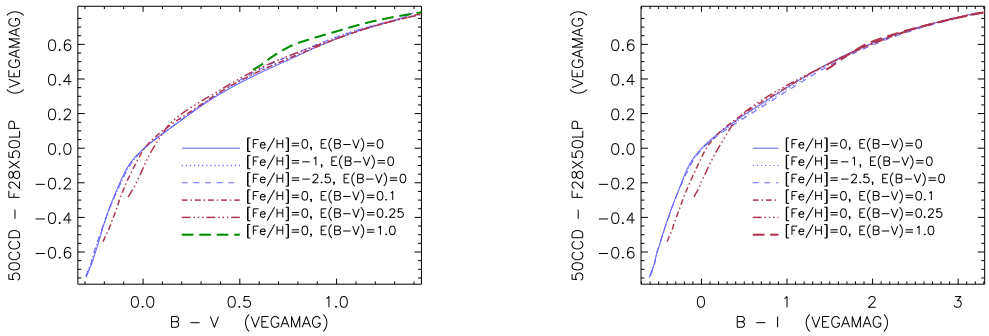

Figure 4: Relations between STIS CCD clear minus longpass color and Johnson $B-V$ and $B-I_c$ colors.

If, for the same series of Kurucz models used for Figure 3, we fit the STIS magnitudes as a sum of the measured SDSS AB fluxes,

$$M_{AB}(STIS) = -2.5 \log \left(\sum_f c_f 10^{-M_{AB}(f)/2.5} \right),$$

we find the best fitting coefficients c_f to be as given in Table 2.

Table 2: Coefficients c_f used to estimate STIS 50CCD and F28X50LP AB magnitudes from SDSS and GALEX magnitudes.

Target STIS		Coefficients of Fit					
Filter	GALEX NUV	u'	g'	r'	i'	z'	RMS residuals
F28X50LP	...	...	...	0.527	0.250	0.224	0.0025
50CCD	...	0.304	0.307	0.131	0.218	0.071	0.019
50CCD	0.070	0.147	0.348	0.209	0.140	0.091	0.002

3.2. Comparison of MAMA and GALEX Passbands

STIS has a variety of broadband UV imaging filters for use with the MAMAs. No particular combination of the available choices was adopted as a standard color, although the FUV MAMA with the SRF2 filter and the NUV MAMA with the QTZ filter were the most commonly used modes. However, the GALEX mission (Martin et al. 2005) will survey the entire sky in NUV and FUV passbands, and this suggests that the GALEX passbands may become such a "standard" vacuum UV color. It is therefore useful to consider how the STIS MAMA imaging modes are related to the GALEX imaging bandpasses. In Figure 5, we compare STIS MAMA and GALEX imaging passbands, and in Figures 5 through 7 we also compare the relation between STIS and GALEX magnitudes as a function of color. Note, however, that the GALEX bandpasses given by Martin et al. 2005 are preliminary, with a quoted accuracy of about 10%.

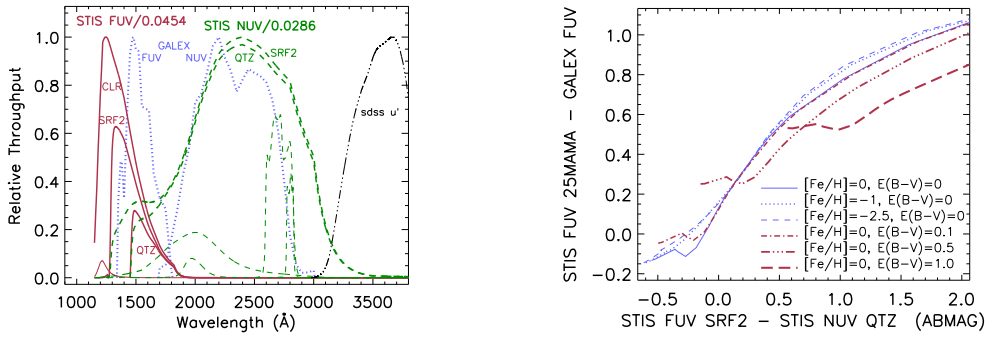

Figure 5: STIS MAMA FUV (solid lines) and NUV (dashed lines) passbands are compared with GALEX passbands (dotted lines) in the left hand panel. On the right, STIS FUV clear aperture and GALEX FUV magnitudes for six series of Kurucz models are compared as a function of a STIS FUV − NUV color.

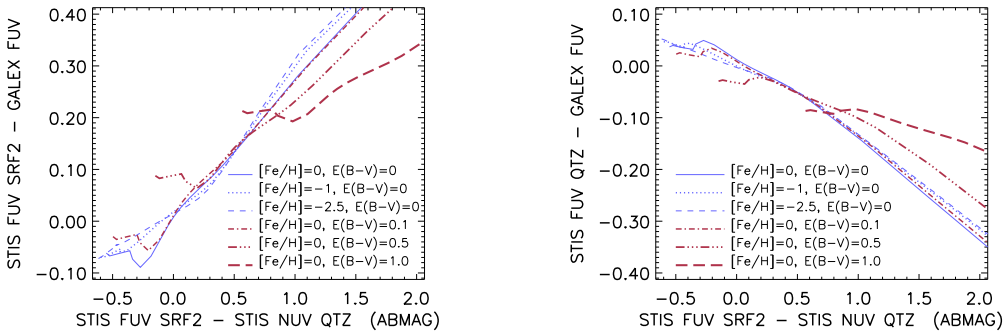

Figure 6: STIS FUV SRF2 (left) and QTZ (right) magnitudes are compared with GALEX FUV magnitudes. Note that the STIS SRF2 filter has a shorter and the STIS QTZ a longer blue cuttoff than does the GALEX FUV bandpass, and this changes the slope of the relationship between the STIS and GALEX magnitudes.

A large fraction of the unfiltered STIS FUV MAMA (25MAMA aperture) throughput is at wavelengths shorter than the blue cutoff of the GALEX FUV filter, and this makes it especially sensitive to differences in the SED. The STIS FUV with the SRF2 and QTZ filters (Figure 6) are better matches. For unreddened stars, relatively tight transformations

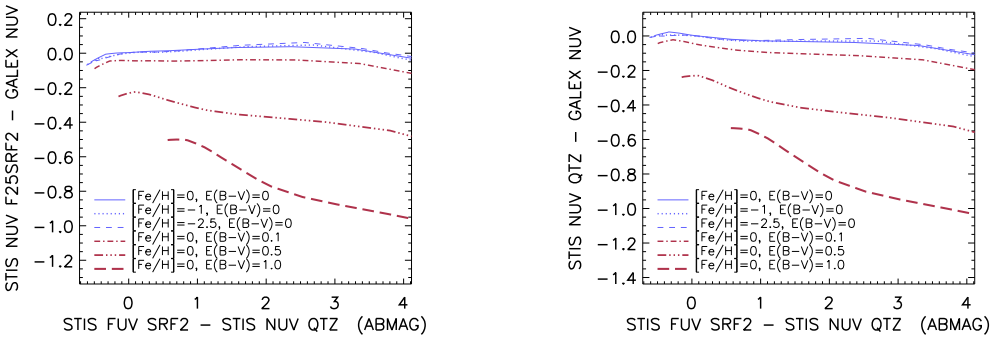

Figure 7: STIS NUV SRF2 (left) and QTZ (right) magnitudes are compared with GALEX NUV magnitudes.

between these systems might be defined, but when extinction becomes substantial, the differences in wavelength cutoffs for these filters have substantial effects.

At both short and long wavelength the throughput of the STIS NUV MAMA QTZ and SRF2 modes extend significantly farther than does the GALEX NUV imaging bandpass. This makes the relationship between STIS and GALEX NUV modes (Figure 7) an especially strong function of the extinction.

4. Summary

In this paper we have only been able to give a brief overview of the STIS imaging mode throughputs and the relationship of these modes to other photometric systems. While many of the details of the calibration are discussed in more detail in Proffitt (2004) and Proffitt et al. (2003), we are also in the process of preparing more detailed quantitative information that will allow STIS imaging mode count rates to be used to approximate measurements in other common photometric bands. Included in this will be formulas that allow time-dependent sensitivity changes, wavelength-dependent aperture corrections, and interstellar extinction corrections to be taken into account.

References

Laidler et al., 2005, *Synphot Data Users Guide* (Baltimore: STScI)

Martin, D. C., et al. 2005, ApJ, 619, 1

Proffitt, C. R. 2004, *Instrument Science Report* STIS 2004-05 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

Proffitt, C. R., Brown, T. M., Mobasher, B., & Davies, J. 2003, *Instrument Science Report* STIS 2003-01 (Baltimore: STScI)

Stys, D. J., Bohlin, R. C., & Goudfrooij, P. 2004, *Instrument Science Report* STIS 2004-04 (Baltimore, STScI)

York, D. G., et al. 2000, AJ, 120, 1579

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Geometric Distortions for the HST MAMA Detectors: STIS NUV-MAMA + FUV-MAMA and ACS SBC

J. Maíz Apellániz[1]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21210, U.S.A.

Abstract. I present the current status of the geometric distortions for the imaging modes of the three current HST MAMA detectors. A solution for the STIS NUV-MAMA with a typical accuracy of 10 mas was implemented in the pipeline two years ago. The same strategy employed there is currently being used to develop new geometric distortion solutions for the STIS FUV-MAMA and the ACS SBC.

1. Description

- We have generated a geometric distortion solution (Maíz Apellániz & Úbeda 2004) for the STIS NUV-MAMA and I am currently working on similar solutions for the STIS FUV-MAMA and ACS SBC, and maybe an improved one for the STIS NUV-MAMA.
- There are two approaches to building a geometric distortion solution:

	Self-calibration	**Astrometric standard field**
Description	Field with multiple orientations and dithering positions but unknown coordinates a priori	Field with one or several orientations and/or dithering positions with known precise coordinates
Advantages	No external data required	Simpler approach
Disadvantages	Many observations required, complex analysis, possible skew issues	Precise standard field needed, more error sources
Examples	Anderson & King (2004), ACS HRC	Maíz Apellániz & Úbeda (2004), STIS NUV-MAMA

- High S/N and knowledge of the PSF are required to obtain accurate positions in the detector frame. S/N is an issue for MAMA detectors due to count rate limitations.
- Few choices for fields of the right brightness, size, and stellar density for MAMA imaging: Galactic globular clusters, scaled OB associations (SOBAs) at 0.8-4.0 Mpc.

2. The current NUV-MAMA geometric distortion

- Astrometric standard fields: NGC 604 (SOBA in M33, $d = 0.84$ Mpc, Freedman et al. 2001) and two in NGC 4214 (galaxy with several massive young clusters, $d = 2.94$ Mpc, Maíz Apellániz et al. 2002, Figure 1).
- Reference coordinates were derived from F170W WFPC2 exposures applying the Holtzmann et al. (1995) solution, which has a typical accuracy of 5 mas (Casertano & Wiggs 2001).

[1]Space Telescope Division, European Space Agency, ESTEC, Noordwijk, Netherlands.

- STIS NUV-MAMA imaging in three filters (CN182, CN270, and SRF2).
- The solution was developed by fitting a fourth-order polynomial degree and appears to be filter independent (Figure 2).
- The typical accuracy is 10 mas.
- The mean scale of the NUV-MAMA is 24.536±0.027 mas/pix in the x direction and 24.795±0.031 mas/pix in the y direction.
- The geometric distortion solution was implemented in the pipeline on 29 October 2003 and is currently in use.

3. New geometric distortion solutions

- Two new astrometric standard fields for the MAMA geometric distortion solutions have been generated from ACS HRC F220W observations. The HRC solution has a typical accuracy better than 0.2 mas, which is more than an order of magnitude better than the Holtzmann et al. (1995) solution for the WFPC2.
- The first one (NGC 604, Figure 3) was already used for the STIS NUV-MAMA. 20 F220W exposures have been obtained under HST program 10722 to generate the standard field.
- The second field (NGC 6681, Figure 4) is a globular cluster at a distance of 9.0 kpc (Harris 1996). 56 F220W exposures obtained under HST programs 9010, 9019, 9565, 9655, 10047, and 10373 were retrieved from the archive to generate the standard field.
- The ACS SBC geometric distortion solution will be generated using 12 positions (each one observed with the F122M, F150LP, and F165LP filters) in the first field using data obtained under HST programs 10419 and 10722; and a total of 230 exposures (using the F122M, F115LP, F125LP, F140LP, F150LP, and F165LP filters) of the second field obtained under ten different HST programs.
- The STIS FUV-MAMA solution will be generated from 110 exposures (using the CLEAR, SRF2, and QTZ filters) of the second field obtained under ten different HST programs.
- If a significant improvement over the previous solution is reached, a new STIS NUV-MAMA solution will be generated from 2 exposures (using the SRF2 and CN182 filters) in the first field obtained under HST program 9096 and 110 exposures (using the CLEAR, SRF2, QTZ, CN182, and CN270 filters) in the second field obtained under ten different HST programs.

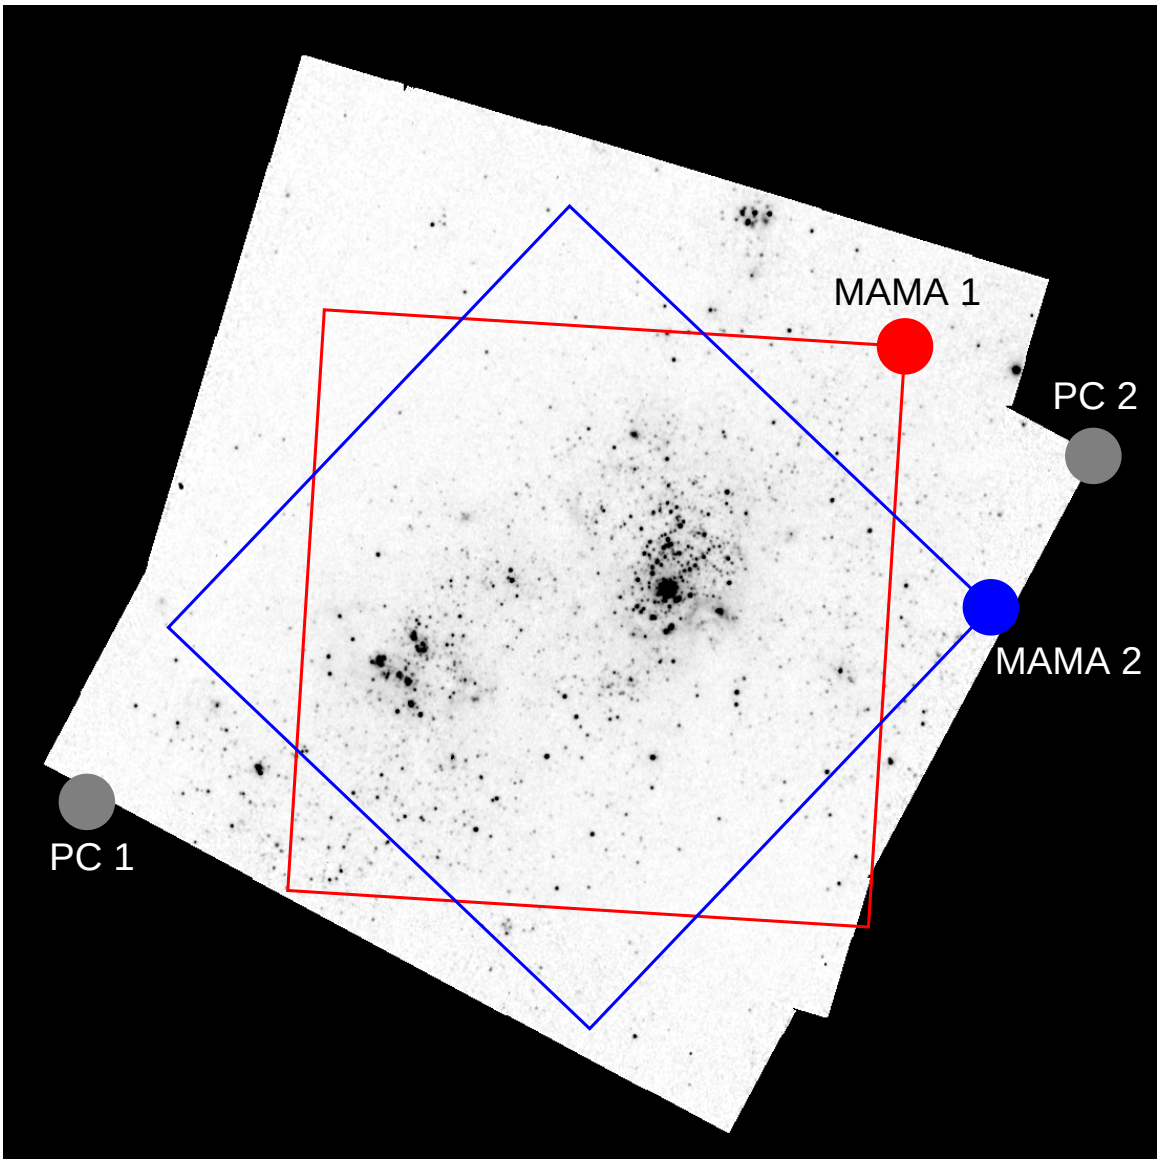

Figure 1: Mosaic with the two F336W PC orientations for one of the NGC 4214 fields with the outline of the two corresponding STIS NUV-MAMA orientations. Circles indicate the lower left corner of the detector for each of the four orientations.

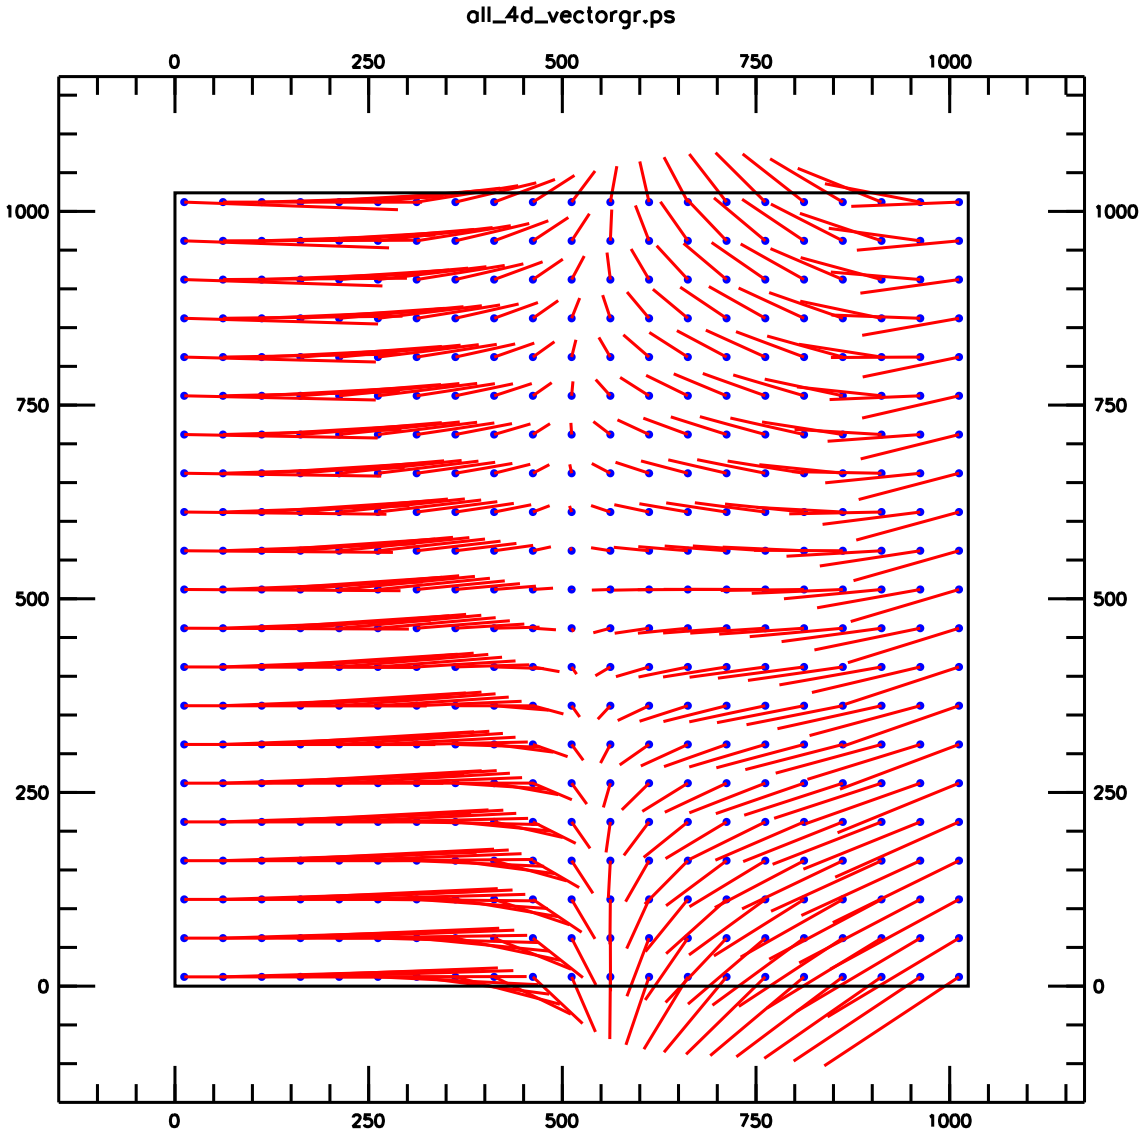

Figure 2: Distortion map for the STIS NUV-MAMA. Each vector represents the distortion in pixels magnified by 50. The dot shows the uncorrected position and the vector the correction to the geometrically correct position.

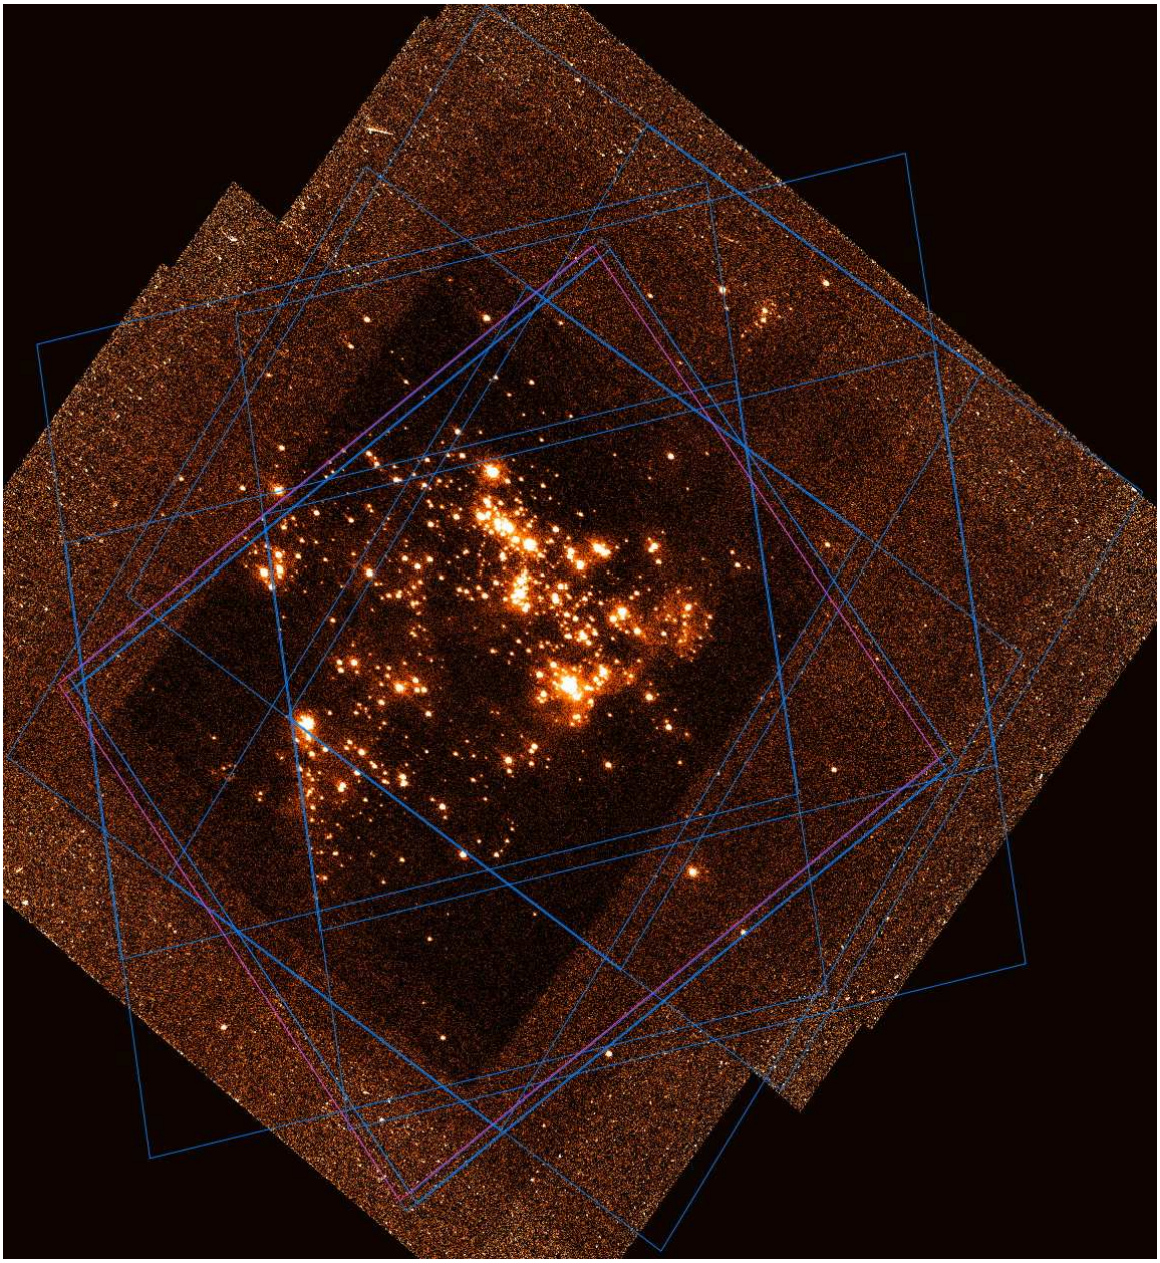

Figure 3: Mosaic with the 20 F220W HRC exposures of the NGC 604 field with the outline of the 12 SBC positions.

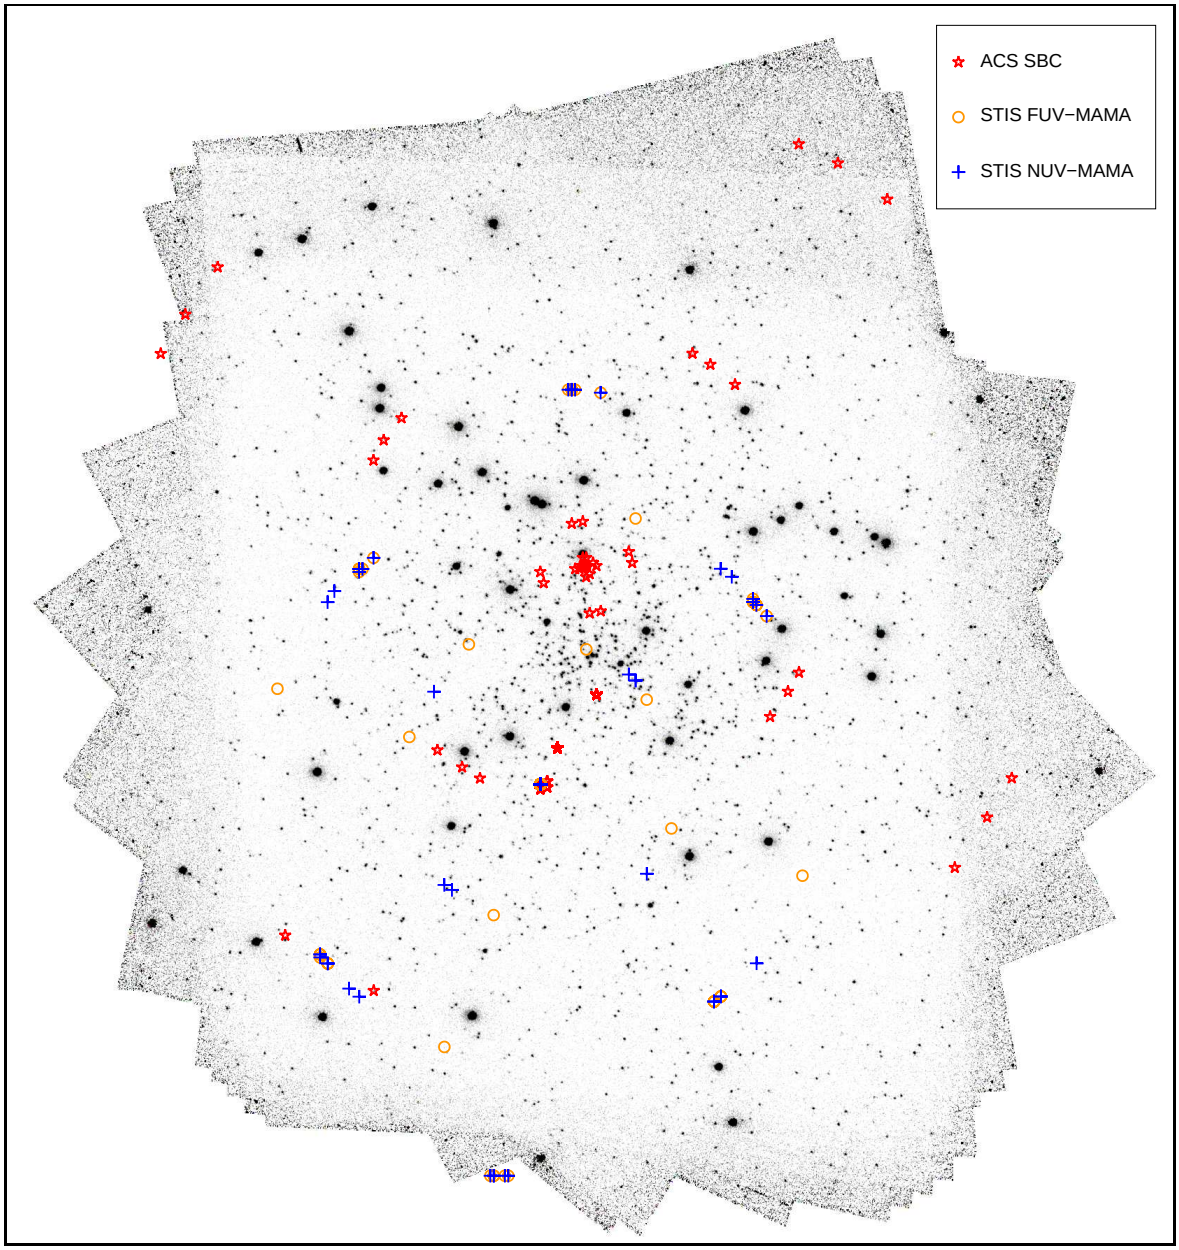

Figure 4: Mosaic with the 56 F220W HRC exposures of the NGC 6681 field. Symbols are used to mark the positions of the MAMA exposures, with many of them corresponding to multiple filters and/or orientations. The field shown is $41'' \times 44''$ with N towards the top and E towards the left. The field sizes for individual MAMA exposures are $35'' \times 31''$ (ACS SBC) and $25'' \times 25''$ (STIS FUV-MAMA and NUV-MAMA).

References

Anderson, J. and King, I. R. 2004, *Instrument Science Report* ACS 2004-15 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Casertano, S. and Wiggs, M. S. 2001, *Instrument Science Report* WFPC2 2001-10 (Baltimore: STScI), available through http://www.stsci.edu/hst/wfpc2

Freedman, W. L., et al. 2001, *ApJ* **553**, 47

Harris, W. E. 1996, *AJ* **112**, 1487

Holtzman, J., et al. 1995, *PASP* **107**, 156

Maíz Apellániz, J., Cieza, L., and MacKenty, J. W. 2002, *AJ* **123**, 1307

Maíz Apellániz, J. and Úbeda, L. 2004, *Instrument Science Report* STIS 2004-01 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Some Neglected Pixel Problems

Kris Davidson
School of Physics and Astronomy, University of Minnesota, 116 Church St. SE, Minneapolis MN 55455

Abstract. We have developed several improved techniques because standard reduction software was inadequate for the Treasury Project on η Carinae. Though adapted specifically to HST/STIS/CCD data, some of these methods can be useful for images obtained with other instruments. For instance, our subpixel modeling technique reduces some of the evil effects of marginal sampling, i.e., in images where the pixels are moderately too large to take full advantage of the instrument's basic point spread function (PSF). Marginal sampling causes the apparent PSF to depend on precise location in the pixel grid, a subtle effect that perturbs measurements. Our method considerably reduces the amplitude of this phenomenon.

1. Preliminary Comment: Not Just Spectroscopy

The techniques sketched here were motivated by intensive STIS spectroscopy of one object. But their usefulness is far more general; they can be applied to *a broad class of images obtained with various instruments.* Whenever the PSF's FWHM is between about 1 and 3 pixel widths, one should consider "subpixel modeling" as outlined in Section 4.

2. The Role of η Carinae Concerning HST Data in General

Before getting into data processing techniques, let me frankly indulge in some advertising which serves a valid purpose here. Most astronomers have a nodding acquaintance with η Car: a very massive star with intense, ostentatiously bipolar ejecta. The degree of structural complexity, however, and the surprisingly basic unsolved problems lurking there, are less well known.[1] Together these characteristics have made Eta unique in HST's record:

- The ejecta are *spatially complex* on size scales of 0.01 to 10 arcsec;
- They and the star are *spectrally complex* too, with thousands of identified emission lines;
- At least six distinct types of spectra are found there, relevant to diverse problems;
- And some of those spectra vary systematically in time.
- Thus we have needed excellent spatial and spectral resolution, appreciable spatial coverage, broad wavelength coverage, and repeated observations, all together.
- Both the star and its ejecta are *bright enough to make such observations feasible* without requiring huge amounts of telescope time.

Scarcely any other major HST target shares this combination of attributes. As a result, η Car has consistently provided unusual illustrations of HST's capabilities. (Why this matters to us here, will become evident in the next paragraph.) Since 1991 this object has held the record for angular resolution in spectroscopy of a complex target, first with

[1] For general information about this object, see Davidson & Humphreys (1997) and three proceedings volumes: Humphreys & Stanek (2005), Gull et al. (2001), and Morse et al. (1999). The HST Treasury Program for Eta has produced numerous scientific results, but here we're concerned with data techniques.

the FOS, then GHRS, and, since 1998, with the STIS. I honestly can't think of another object that demonstrated the STIS as well – for instance, almost nothing else was observed across the entire wavelength range, as we did repeatedly for Eta. Several HST instruments have been used *intensively* on this object, close to their resolution limits, sometimes with unconventional techniques. The 72-orbit Treasury Project obtained a vast number of meaningful data pixels (more than a typical multi-hundred-orbit extragalactic program, for instance), because many integrations could be obtained in each orbit and because of the extended structure.[2] In some respects this was the most ambitious spectroscopic program ever attempted with the HST.

Standard reduction software and standard information about the instrument proved to be inadequate for such intensive observations. Therefore, entirely apart from scientific goals, *these data were and are valuable in three practical ways:* (1) They clearly and quantitatively revealed various glitches in the STIS/CCD and in the data reduction procedures; (2) they forced us to develop techniques to reduce those problems; and (3) they provide a unique, public archival record of instrument characteristics as functions of time. The entire set of STIS data on η Car, including pre-project observations incorporated into the overall data set, extend from 1998.0 to 2004.2, most of the instrument's working life.[3]

We hope these facts may be useful for improved analyses of existing STIS data on other objects – particularly those where the spatial resolution proved to be unsatisfying because of shortcomings in the reduction software, not the instrument itself (see below). Some of the techniques can be applied to other instruments as mentioned earlier.

3. A Brief List of Glitches

Thus, as foreseen in the original proposal, a substantial fraction of our effort in the η Car Treasury Program has gone into improving the STIS/CCD data processing methods and assembling a convenient, somewhat unusual data archive. In this section I'll just list the main difficulties we've dealt with; here "we" means mainly John Martin, Kazunori Ishibashi, and myself, the persons working on these development tasks.[4] Additional information on the following effects and techniques can be found at the project website, `http://etacar.umn.edu`.

(1) At UV to visual wavelengths the STIS CCD provided only marginal spatial sampling – i.e., the 0.05 arcsec pixel size was undesirably large relative to the PSF width. With standard software, *this fact severely degraded the spatial resolution of spectroscopy.* The reason, which concerns variations in local sampling, is more complicated than one might assume at first glance. In order to get resolution better than 0.2 arcsec with this instrument, one must extract a spectrum less than 4 CCD rows wide. Since a spectral "trace" is slightly tilted relative to a CCD row, the spatial sampling varies as a function of wavelength; at some detector columns the spectrum of a point source coincides with a row center, while at others it falls midway between two rows. With normal interpolation methods, this systematic variation in sampling causes any narrow continuum extraction to appear "scalloped" or wavy. If one attempts to observe a target near another object, the contamination cannot accurately be subtracted out – at least, not with standard software – because its undulations

[2] Of course 72 orbits is a big number by most users' standards. But those orbits were required to study intricate variations applicable to several branches of astrophysics, not just a basic set of spectra.

[3] One potentially useful detail: Certain ejecta blobs have rich spectra of narrow, well-classified emission lines, which provide supplementary wavelength calibrations.

[4] For some purposes Matt Gray, Michael Koppelman, J. T. Olds, and R. M. Humphreys should also be included in the list. Other project CoI's have been concerned mainly with various aspects of the scientific results, as planned.

generally differ in both phase and amplitude from those of the target object. We suspect that this phenomenon discouraged some early STIS users from making further attempts to do high-spatial-resolution spectroscopy. Dithering could have improved the sampling but was not always feasible – for instance, we didn't have enough telescope time to dither the Treasury Project spectroscopy. Fortunately, a technique sketched in Section 4 below – essentially a novel form of pixel interpolation – provides a considerable improvement. This method can be applied to existing STIS data and to other image data whose pixels are somewhat larger than one would have preferred.

(2) Because our subpixel modeling technique (effect 1 above and Section 4 below) depends on local arrays of about 20 pixels, it was necessary to eliminate practically all bad pixels – hot pixels, cosmic ray hits, etc. Therefore we employed a bad-pixel identification method in addition to the normal CR-SPLIT procedure. To test each pixel, our software analyzed the dispersion of pixel values in a 5×5 sample array centered at that location. This technique works rather well even if only one image is available, i.e., if no CR-SPLIT companion image exists.

(3) The complete STIS/CCD wavelength range, UV to far red, required about 30 grating tilts. In order to produce a unified spectral image covering the entire range, we need to splice each adjoining pair of images. However, there is an unexpected complication: the detector was not uniformly focussed. In fact the PSF was generally sharper at the longer-wavelength side of each image. Therefore we have had to develop a convolution procedure to make adjacent wavelength samples match in their overlap region. This is further complicated because tiny variations in pointing caused the apparent star/ejecta brightness ratio to vary somewhat.

(4) Each STIS/CCD spectral image contains several types of "ghosts" with various intensities and size scales. We have parametrized the worst of these, especially those with small size scales. With this information they can be reduced by a simple procedure that basically amounts to Jansson – Van Cittert deconvolution.

(5) The STIS/CCD spatial PSF had asymmetric wings. We have measured the parameters necessary to reduce this effect, in much the same manner as effect 4 above.

(6) Suitable ERR arrays – local noise levels stored in each FITS file – are not entirely obvious. In particular, when the stored pixel values have been produced by interpolation between original data pixels (see Section 4 below), then one should *not* store the mathematically correct r.m.s. error for each interpolated pixel. Reason: The interpolation process causes adjacent pixel values to be correlated with each other. In that case one cannot estimate the true error of a sum of pixels merely by taking a quadratic sum of the individual formal errors. – I recognize that these comments may seem a little cryptic, but a full explanation would take too much space here so the reader is invited to examine the problem himself/herself. The point is that the ERR arrays in some pipeline-processed STIS images are misleading, and don't contain enough information to reconstruct the true error in a sum or average of many pixels. In fact an uncritical application of them will usually lead to an underestimate of the noise level, sometimes a serious underestimate.

There are other difficulties but these are probably the most troublesome. Items 1, 2, and 6 apply to image data in general, and effect 3 might also have applications beyond STIS. The remainder of this paper concerns item 1.

4. Optimum Subpixel Modeling

For quantitative work we usually need to correct optical distortions, to rotate the image, to make position measurements, etc. These tasks require some form of interpolation between the original pixel locations. Often this "subpixel modeling" is implicit or hidden in the software, but it is essential; and marginal sampling makes it trickier.

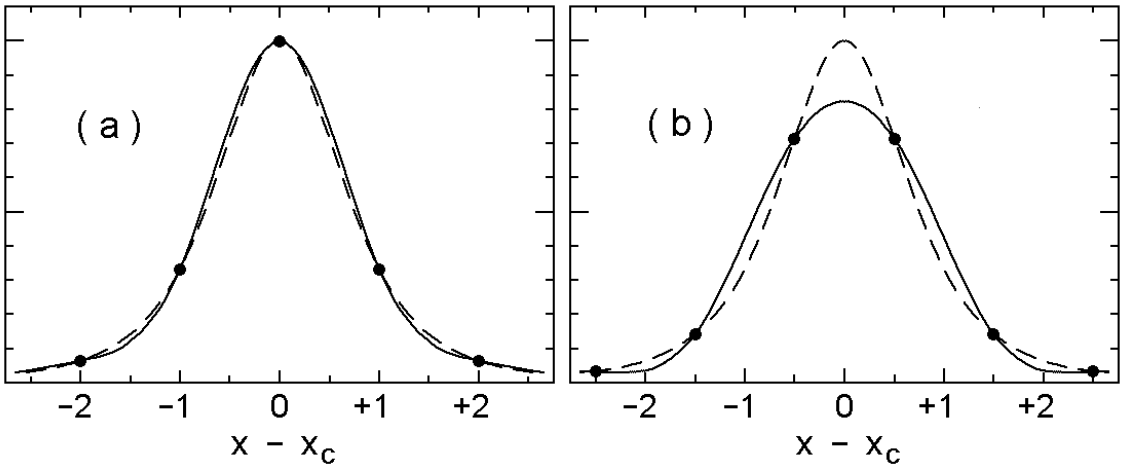

Figure 1: The dashed curves represent a function with FWHM = 1.5 pixels. The continuous curves are spline fits to the sample points shown as dots. Here $f(x)$ is the square of a Lorentzian, $[\,1+\alpha(x-x_c)^2\,]^{-2}$, which closely resembles the PSFs found in STIS/CCD data.

By "marginal sampling" I mean any case where the PSF width (FWHM) is between, roughly, 1 and 3 pixel widths. Sampling naturally degrades the average resolution, but that isn't the main problem. Explanations are easier if we idealize the situation in two ways: (1) Rather than a two-dimensional image $f(x, y)$, let's consider the one-dimensional case $f(x)$. Our unit of x will be the original pixel width. (2) In a real image each pixel value is some sort of average across the pixel, but it's easier to discuss uniformly spaced samples of a continuous function, $f_i = f(x_i)$. Neither of these simplifications affects the basic problem; the approach described below can be generalized to real data.

Figure 1 shows a function – maybe a PSF – sampled with two different pixel offsets. In Fig. 1a the peak coincides with a data point (a pixel center), and a cubic spline through the data points matches the function pretty well. In Fig. 1b, however, the peak falls midway between two sample points; then the spline fit is much worse, with a severely depressed peak and a broadened profile. Evidently *the effective PSF depends on a point source's location on the detector.* This fact is fundamental, not merely a consequence of the interpolation scheme. The sampling is blind to some information near and beyond the critical Fourier frequency $\omega_c = \pi$ pixel^{-1}, and *the precise location of the pixel grid determines which Fourier components are lost.* One doesn't notice the resulting spatial fluctuation of the PSF when viewing an image, but it affects measurements. This phenomenon is responsible for the undulations in narrow spectral extractions, effect 1 in Section 3 above. Incidentally, one can show that deconvolution methods don't help this problem much.

In principle the effective sampling frequency can be improved by "dithering," taking multiple exposures with fractional-pixel relative offsets. But dithering requires extra effort and can be inefficient in telescope usage; it depends on reliable, extremely precise offset pointing; it can fail for various reasons without the failure being obvious; and sometimes one simply doesn't have enough telescope time. A *consistent* resolution of about 2 pixels (e.g., 0.1 arcsec with the STIS/CCD) can be attained without dithering, if we use a special technique outlined below. This subpixel modeling technique greatly reduces the spatial variation of the PSF.

Here's our plan. We use the original data points to generate a related function $g(x)$ with the following goals:

1. $g(x)$ approximates $f(x)$ as closely as possible without excessively violating requirement 2 below. Function $g(x)$ will turn out to be a blurred version of $f(x)$ and of course we hope to minimize the blurring.

2. A second goal, motivated by Fig. 1, distinguishes our effort from ordinary interpolation: *The calculated shape of $g(x)$ must depend as little as possible on the location of $f(x)$ relative to the pixel grid.* For instance, imagine two images $f_1(x)$ and $f_2(x)$ with the same shape but relatively offset by half a pixel: $f_2(x) = f_1(x - 0.5)$. We design the corresponding g-functions so that $g_2(x) \approx g_1(x - 0.5)$ as accurately as feasible. As Fig. 1 implies, this requirement is subtle. It isn't easy to achieve without heavy blurring.

3. Based on $f(x_i)$ at integer points $x_i = i$, we calculate $g(x_k)$ with half-integer spacing, $x_k = k/2$. Since this reduced interval provides decent sampling, we can later estimate $g(x)$ at any other point by conventional interpolation.

Goals 1 and 2 tend to conflict with each other, so the "best" solution for any given data set is a compromise. Moreover, since quantitative criteria depend on the specific application, there can be no universal solution. Our reasoning will be fairly general.

Since the procedure for generating $g(x)$ must be linear in $f(x)$, each output sample point $g(x_k)$ must be a linear combination of nearby input data $f(x_i)$. We further assume symmetry between the positive and negative x-directions; this is not strictly true for some instruments (including STIS) but it's a good practical approximation that greatly simplifies the effort. These two requirements determine the formula for locations that coincide with the original data points x_i. If x_k is an integer, then

$$g(x_k) \;=\; A_0\, f(x_k) \;+\; \sum_{n=1}^{M-1} A_n \cdot \{\, f(x_k - n) + f(x_k + n) \,\} \;, \tag{1}$$

where the A_n are constant coefficients. Similarly, for the intermediate output points where $x_k =$ integer + 1/2,

$$g(x_k) \;=\; \sum_{n=0}^{N-1} B_n \cdot \{\, f(x_k - 0.5 - n) + f(x_k + 0.5 + n) \,\} \;. \tag{2}$$

In a practical scheme M and N will usually be 2, 3, or 4. "Conservation of counts" requires

$$A_0 \;+\; 2 \sum_{n=1}^{M-1} A_n \;=\; 2 \sum_{n=0}^{N-1} B_n \;=\; 1 \;. \tag{3}$$

Moreover, if

$$\sum_{n=1}^{M-1} n^2\, A_n \;=\; \sum_{n=0}^{N-1} (2n+1)^2\, B_n \;=\; 0 \;, \tag{4}$$

then $g(x)$ exactly matches $f(x)$ if the latter is any cubic polynomial. Constraint 3 is essential but 4 is not.

Let me emphasize that *any* linear symmetric scheme will be equivalent to eqns. 1 and 2, even if it's presented in terms of Fourier analysis or other math language.[5] $|A_{n+1}|$ tends to be much smaller than $|A_n|$ and formula 1 either blurs or sharpens $f(x)$, depending on

[5]If we drop the assumption of symmetry in the x-direction, then we still have linear combinations of $f(x_i)$, but with almost twice as many coefficients.

whether A_1 is positive or negative. For reasons that I don't have space to prove here, we must accept the former, the blurred case. Formula 2 is like standard interpolation but the coefficients are based on non-standard criteria. If we later generalize to realistic pixels, i.e., local averages of $f(x)$, only the coefficient values will be altered.

The coefficients A_n and B_n will be based on goals 1 and 2 stated earlier. A criterion for goal 2 can be defined as follows. First, choose a test function $f(z)$ like that shown in Fig. 1; here we'll use coordinate $z = x - x_c$ instead of x, where x_c is some arbitrary offset. The test function peaks at $z = 0$, it's normalized so that $f(0) = 1$, and it resembles the instrumental PSF. Consider any particular value of z. If we sample the function at z, $z \pm 1$, $z \pm 2$, etc., then a g-value can be calculated from eqn. 1:

$$g_A(z) \;=\; A_0\, f(z) \;+\; A_1 \cdot \{f(z-1) + f(z+1)\} \;+\; ...$$

If, instead, we sample it at $z \pm 0.5$, $z \pm 1.5$, etc., then eqn. 2 becomes applicable:

$$g_B(z) \;=\; B_0 \cdot \{f(z-0.5) + f(z+0.5)\} \;+\; B_1 \cdot \{f(z-1.5) + f(z+1.5)\} \;+\; ...$$

The difference $g_B(z) - g_A(z)$ indicates the local departure from requirement 2. A reasonable measure of the overall discrepancy is

$$E_{AB} \;=\; \Big[\int_{-\infty}^{+\infty} \{g_B(z) - g_A(z)\}^2 \; dz \Big]^{1/2} \; . \tag{5}$$

We adjust the coefficient sets A_n and B_n to minimize E_{AB}, subject to other requirements. Parameter space must be explored with many calculations – at least I haven't thought of any other way to do it. For some applications, it may be useful to emphasize selected parts of the PSF by means of a weighting function $w(z)$ in the above integral.

Fourier considerations show why there is no really satisfying set of coefficients. If the test function $f(z)$ is symmetric, then we need only the cosine transforms $F(\omega)$, $G_A(\omega)$, and $G_B(\omega)$, referring to $f(z)$, $g_A(z)$, and $g_B(z)$. Then formulae 1 and 2 have the following effects:

$$G_A(\omega) \;=\; H_A(\omega)\, F(\omega) \qquad \text{and} \qquad G_B(\omega) \;=\; H_B(\omega)\, F(\omega) \; ,$$

where

$$H_A(\omega) \;=\; A_0 \;+\; 2A_1 \cos\omega \;+\; 2A_2 \cos 2\omega \;+\; ... \; ,$$

$$H_B(\omega) \;=\; 2B_0 \cos 0.5\omega \;+\; 2B_1 \cos 1.5\omega \;+\; ...$$

$H_A(\omega)$ and $H_B(\omega)$ are filtering functions that dampen Fourier components near the critical frequency $\omega_c = \pi$, and the parameter E_{AB} (eqn. 5) can be expressed in terms of a frequency integral of $\{\,(H_B - H_A)\,F\;\}^2$. Our "goal 1" stated earlier requires a rapid cutoff just below ω_c, while goal 2 requires $H_A(\omega) \approx H_B(\omega)$. Unfortunately, as Fig. 2 (next page) shows, *the two filtering functions are largely incompatible at high frequencies* – because $dH_A/d\omega = 0$ and $H_B = 0$ at $\omega = \omega_c$, two very different conditions. Moreover, H_A and H_B have opposite slopes above the critical frequency. Since these are fundamental consequences of the sampling, goals 1 and 2 cannot both be satisfied to high accuracy.

Finding a suitable compromise is difficult. We can force H_A to H_B resemble each other for $\omega < \omega_c$ by imposing the additional constraints $H_A = 0$ and $dH_B/d\omega = 0$ at ω_c. Then, however, we find that the cutoff occurs at lower-than-desired frequencies, which means that the test function $f(z)$ is severely blurred. In order to lessen the blurring, we must accept a set of coefficients that causes H_A to dip slightly below zero at ω_c where H_B automatically passes through zero. This gives acceptable results provided that the cosine transform $F(\omega)$ is small at the critical frequency – which is a fancy way of saying that $f(z)$ must not be too narrow.

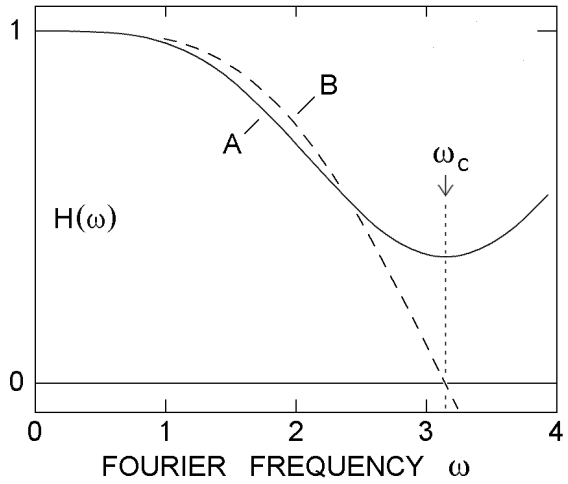

Figure 2: General shape of the filtering functions $H_A(\omega)$ and $H_B(\omega)$ for typical sets of coefficients. At the critical frequency, H_A is level while H_B passes through zero. (The particular coefficients used here do not give a useful solution to our problem, but were chosen merely to illustrate behavior near ω_c.)

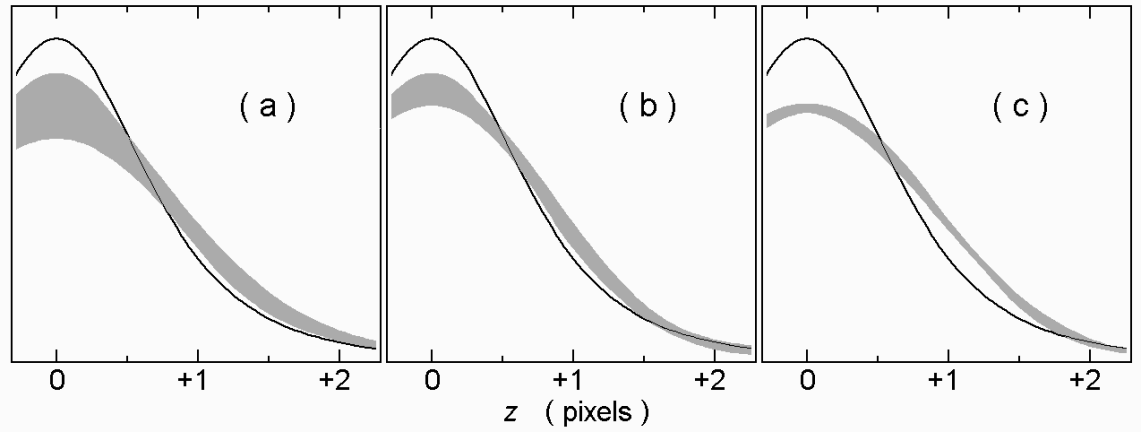

Figure 3: Interpolation envelopes for (a) linear interpolation, (b) spline interpolation, and (c) eqns. 1 and 2 with coefficients mentioned in the text. In each case the solid curve represents the instrumental PSF before it is sampled by the detector. After sampling and interpolation, the effective PSF depends on the precise subpixel location of the point source (compare Fig. 1); gray areas here represent the envelopes of interpolated values for all possible pixel-grid offsets.

In fact the necessary blurring doesn't hurt as much as one might imagine. Referring back to Fig. 1, the effective output PSF will be a little broader than the solid curve of case 1b. This *appears* much worse than 1a. but the resolution shown in Fig. 1a is largely illusory; it depends on placing the pixel grid in an advantageous position and then knowing that this has been done. Fig. 1b gives a more realistic impression of the true resolution for most applications.

Here I won't recommend specific A_n, B_n coefficient values, for three reasons: Several alternatives would need to be listed because the criteria depend on instrument parameters; frankly I'm not yet sure which choices work best in general; and, anyway, this isn't a reference work. One very simple example is worth noting, though. Early in our Treasury Project work, before being fully aware of all the above considerations, we chose the following coefficients for STIS/CCD data: $A_0 = +0.77$, $A_1 = +0.165$, $A_2 = -0.05$, $B_0 = +0.58$, and $B_1 = -0.08$. These were applied to real pixel values, not function samples $f(x_i)$. Fig. 3 shows results compared to linear and spline interpolation used in most software. For any method, the effective output PSF depends on the point source's precise subpixel location as explained earlier. The solid curve in the figure represents a pre-detector PSF; the other curves in the figure are broader and lower because they represent averages over realistic pixels. Each shaded area is an envelope of profiles that occur for all possible subpixel locations.

As Fig. 3 demonstrates, our method (with the coefficients noted above) produces a substantially tighter envelope than either linear or spline interpolation. On the other hand it is not as good as we would like. This figure portrays the behavior of blue-wavelength STIS data quite well, and here I've neglected some details – such as wriggles in the profile wings – too complicated to discuss now. Increasing the number of coefficients produces a modest, not dramatic, improvement in performance.

For more information see `http://etacar.umn.edu`.

5. Acknowledgments

The Hubble Treasury Project for η Car is of course supported by funding from STScI. Some of the work reported here was done by John Martin at Minnesota and Kazunori Ishibashi at MIT, with considerable assistance from Matt Gray, Michael Koppelman, J.T. Olds, and R.M. Humphreys. Martin, Ishibashi, and I are grateful to Paul Goudfrooij, Linda Dressel, and Paul Barrett for valuable and interesting discussions of these topics.

References

Davidson, K., & Humphreys, R.M. 1997, ARA&A, 35, 1

Gull, T.R., Johansson, S., & Davidson, K. (eds.) 2001, Eta Carinae and Other Mysterious Stars, ASP Conf. Ser. 242

Humphreys, R.M. & Stanek, K.Z. (eds.) 2005, The Fate of the Most Massive Stars, ASP Conf. Ser. 332

Morse, J.A., Humphreys, R.M., & Damineli, A. (eds.) 1999, Eta Carinae at the Millennium, ASP Conf. Ser. 179

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The η Carinae Treasury Project and the HST/STIS

John C. Martin and Kris Davidson

School of Astronomy & Physics, University of Minnesota, Minneapolis, 55455

The Eta Car Treasury Project Team

listed in the Acknowledgments

Abstract. The HST Eta Carinae Treasury Project made extensive use of the HST/STIS from 1998 to the time of its failure in 2004. As one of the most prolific users of that instrument, the Treasury Project used the cross-dispersed spatial resolution of the STIS as few projects did. We present several enhancements to the existing STIS data reduction methods that are applicable to non-Treasury Project data in the STIS archive.

1. Introduction

η Car is an unusually demanding target for HST. The standard software is inadequate for several reasons. We mentioned some of these problems in an article in Davidson (2004a), and you can find details in our Technical Memos at our website[1]. These items specifically concern STIS/CCD, but some have broader applications. STIS users should be aware that the Tech Memos on our site are updated from time to time.

2. Oversized Pixels

HST instruments generally have marginal or poor spatial sampling – e.g., each STIS/CCD pixel was about 50% wider than it should have been for the optical system. Dithering can improve the sampling but is cumbersome and is not always feasible. We have devised a special interpolation method to minimize the bad effects that occur when one applies distortion corrections, rotations, etc., to poorly sampled data. Figure 1 indicates the improvement in a spectrum extraction, for example, but this technique can and should be applied to all non-dithered STIS, WFPC2(PC), and ACS data. See the paper by Davidson (2006) in this volume and our Technical Memo 1 (Davidson 2004b) for more details.

3. Subtleties Concerning Noise

Many astronomers use the ERR arrays in HST data files, but this can be dangerous. For instance, Figure 2 shows the ERR array in a pipeline-reduced STIS/CCD spectral image. The beautiful pattern does not represent any fundamental instrument effect; it is a processing artifact related to pixel resampling. The cause is too subtle to explain here, but for most practical purposes the standard ERR array is wrong and one cannot estimate the true statistical noise from it. Our special interpolation technique mentioned above avoids

[1] `http://etacar.umn.edu/treasury/techmemos`

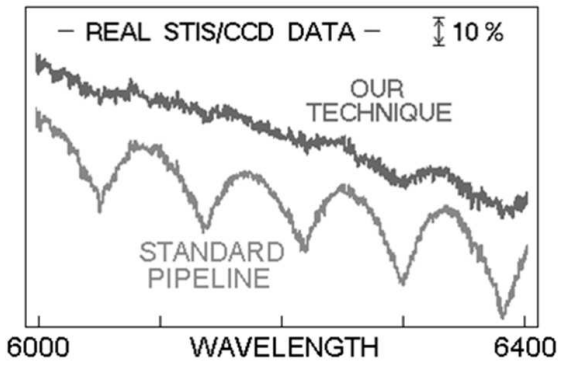

Figure 1: Slight off-center stellar continuum, extraction with a width of $0.1'' = 2$ STIS/CCD pixels.

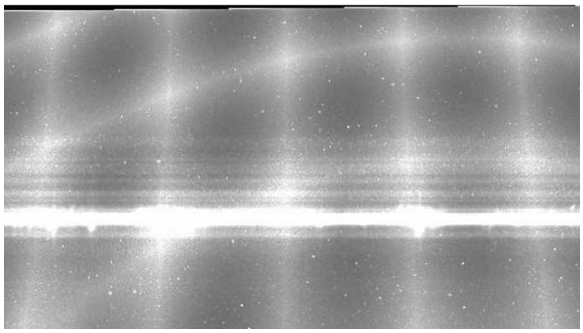

Figure 2: The ERR array in pipeline-processed STIS data, see Section 3.. There is a mathematical reason for the pattern, but for most purposes it gives wrong answers.

this pitfall, and the ERR arrays in our archived Treasury Project data will be carefully and unambiguously defined. All this is explained in Treasury Project Technical Memos 1 and 7 (Davidson 2004a, 2004c), which apply to any data that have been resampled, corrected for distortions, etc.

4. Low-level Wings of the STIS/CCD Spatial Point Spread Function (PSF)

The left part of Figure 3 shows three typical log-scale plots of the weak but extensive response wings along the STIS slit. Since these matter in the case of η Car, we have assessed the parameters and developed a correction routine. The right-hand part of Figure 3 shows the before and after affect of this correction in a real STIS spectrogram of an isolated star. Note that some complex structure remains, which we have not yet quantified. See our Technical Memo 6 (Martin 2005a).

5. Asymmetry in the Core of the PSF

At long wavelengths the STIS/CCD spatial PSF has a very noticeable bump on one side, close to the peak and much stronger than the wings mentioned above (Figure 4). We have parametrized this effect and we will soon produce a simple routine to remove it if a user so desires. See Technical Memo 2 (Martin 2004).

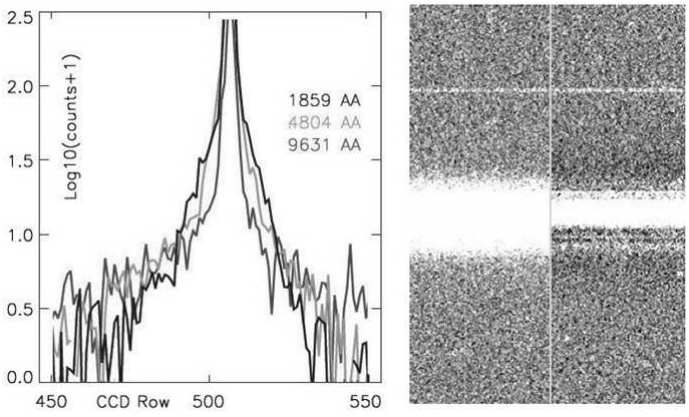

Figure 3: Left panel: the level of extended wings from a point source at different wavelengths. The peak of the PSF core is at about $\log(10^4) = 4$. Right panel: an example of the before and after effect of removing these wings.

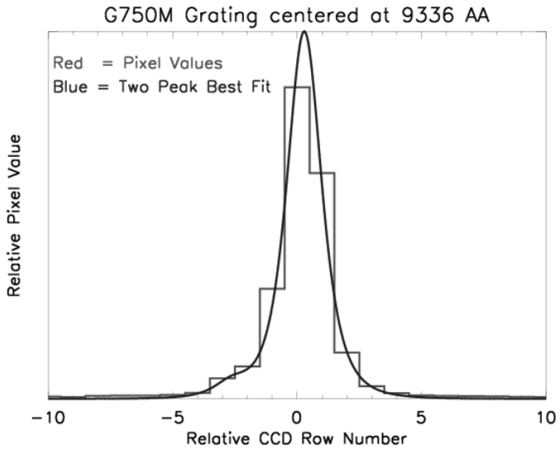

Figure 4: A cross cut of a point source spectrum on the STIS/CCD showing the asymmetry in the PSF. The flux scale is linear so the wings from Fig 3 are not discernible.

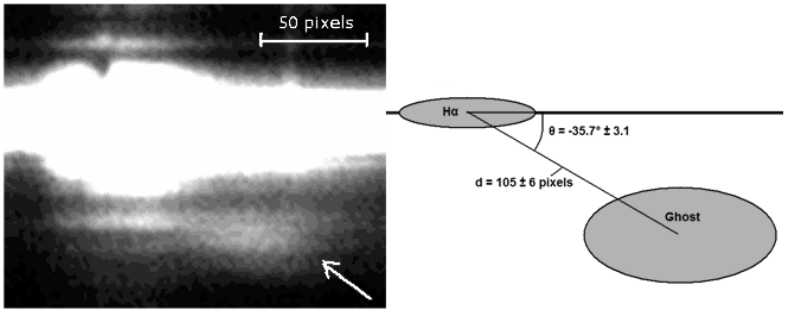

Figure 5: A demonstration (left) and diagram (right) of the H-alpha ghost.

6. The "H-alpha Ghost"

Since the star's Ha emission line is extremely bright, scattering or reflection within the instrument produced an obvious, extended ghost image nearby (Figure 5). To a useful extent we can remove this ghost in order to clarify faint underlying nebulosity. This problem is described in Technical Memo 10 (Martin 2005c).

7. Bad pixels

We supplement the familiar "CR split" method for identifying cosmic ray hits with a second technique that does not require two independent exposures. This method will be described in another Technical Memo (not yet available).

8. Saturated pixels

In the Hα emission line and in a few long integrations, the Treasury Project data contain many saturated pixels. In most such cases we obtained shorter integrations which allow patching. The archive will include both patched and unpatched versions. See Technical Memo 8 (Martin 2005b).

9. Splicing data taken with successive grating tilts

Naturally we wish to produce unified two-dimensional spectral images spanning the entire CCD wavelength range from 1700 Å to 10000 Å. Unfortunately the instrument focus varies across the CCD. Therefore we must employ special convolution techniques to make two adjoining spectral samples (using successive grating tilts) match in their overlap. This will be completed by Spring 2006.

10. Other effects

The above is by no means a complete list! We are working on other details and structural problems.

Acknowledgments. The P.I. of the HST Eta Carinae Treasury Project is K. Davidson (UMN). Observing plans were prepared by T.R. Gull (GSFC). K. Ishibashi (MIT) modified the initial-processing software and performed the first-stage data reduction. Subsequent development and analysis work reported here have been done at UMN mainly by J.C. Martin and K. Davidson, assisted by M. Gray, R.M. Humphreys, and M. Koppelman.

Other proposal CoIs at various institutions (alphabetically): M. Bautista, M. Corcoran, A. Damineli, F. Hamann, H. Hartman, D.J. Hillier, S. Johansson, O. Stahl, N.R. Walborn, and K. Weis.

This work is supported by grants GO-9420 and GO-9973 from the Space Telescope Science Institute, which is operated by AURA and supported by NASA.

References

Davidson, K. 2004a, STScI Newsletter, Spring 2004, 1

Davidson, K. 2004b, "Our Adopted Scheme for Subpixel Modeling," η Car Treasury Project Tech Memo, 1 [2]

Davidson, K. 2004c, "The Meaning of ERR HDU's or Noise Files," η Car Treasury Project Tech Memo, 7 [3]

Davidson, K. 2006, this volume, 247

Martin, J.C. 2004, "The STIS/CCD Cross-Dispersion Point Spread Function," η Car Treasury Project Tech Memo, 2 [4]

Martin, J.C. 2005a, "A Model for the Distribution of Scattered/Diffuse Light on the STIS CCD in Spectroscopic Mode," η Car Treasury Project Tech Memo, 6 [5]

Martin, J.C. 2005b, "Identifying and Patching Overexposed Pixels in STIS CCD Data," η Car Treasury Project Tech Memo, 8 [6]

Martin, J.C. 2005c, "Removal of the H-alpha Ghost Image From STIS CCD Data," η Car Treasury Project Tech Memo, 10 [7]

[2] `http://etacar.umn.edu/treasury/techmemos/pdf/tmemo001.pdf`

[3] `http://etacar.umn.edu/treasury/techmemos/pdf/tmemo007.pdf`

[4] `http://etacar.umn.edu/treasury/techmemos/pdf/tmemo002.pdf`

[5] `http://etacar.umn.edu/treasury/techmemos/pdf/tmemo006.pdf`

[6] `http://etacar.umn.edu/treasury/techmemos/pdf/tmemo008.pdf`

[7] `http://etacar.umn.edu/treasury/techmemos/pdf/tmemo010.pdf`

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Spectral Extraction of Extended Sources Using Wavelet Interpolation

Paul Barrett

Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218

Linda L. Dressel

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. Spatially-resolved spectroscopic investigations of extended sources using STIS, such as galactic black hole mass measurements and chemistry in spatially resolved structures around Eta Carinae, use a narrow extraction width (1-2 pixels wide) to minimize source confusion. However, the combination of the finite pixel size of the CCD and the slight tilt of the spectrum cause problems for the CALSTIS extraction algorithm, resulting in spectra having a scalloped pattern (i.e. aliasing). A wavelet interpolation algorithm is presented that nearly eliminates aliasing, while preserving the flux content of the subpixels. This algorithm will be implemented as a STSDAS task.

1. Introduction

Spatially-resolved spectroscopic investigations of extended sources using the Space Telescope Imaging Spectrograph (STIS), such as galactic black hole mass measurements (Dressel 2003) and chemistry in spatially resolved structures around Eta Carinae (Davidson 2004), use a narrow extraction mask (1-2 pixels wide) to minimize source confusion. However, the narrow mask causes the spectra to show a scalloped pattern (Figure 1). This strong ($\sim 15\%$) aliasing is due to the point spread function (PSF) being only marginally sampled (FWHM ~ 1.3 pixels) and the spectral trace (i.e. the peak or ridge of the spectrum) being slightly tilted with respect to the rows of the detector (Figure 2). For point-source spectra, a mask with a large (7 pixels for CCD modes) extraction width is used, which encompasses all of the flux and avoids the aliasing problem. To minimize aliasing, the STIS Instrument Handbook (Kim Quijano et al. 2003) recommends dithering spectral observations along the slit by a half pixel. This is often not done, because it doubles the requested exposure time, perhaps resulting in a less competitive proposal. When done, it reduces the magnitude of the aliasing, but also doubles its frequency, which can be problematic for spectral features of comparable scale. This paper describes a wavelet interpolation algorithm that nearly eliminates the aliasing problem for narrow extraction masks, while preserving the flux content of the interpolated spectral image.

2. Average Interpolation

Interpolation is widely used in the calibration and analysis of astronomical data to transform images to a common frame or uniform grid. The most commonly used algorithms are linear and bi-linear interpolation for one- and two-dimensional data, respectively. These simple algorithms are usually sufficient for a rough estimate of the interpolated value, but can give erroneous results when high precision is needed. This is particularly true when the

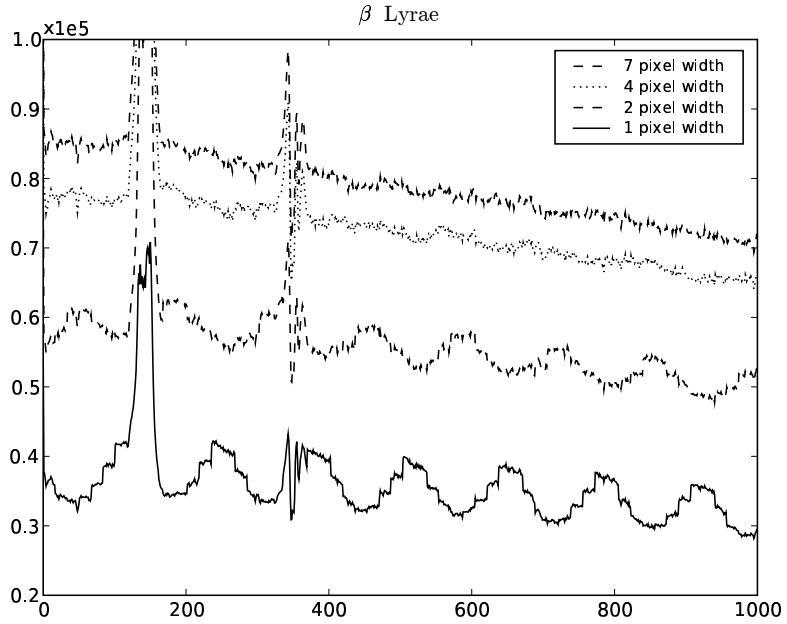

Figure 1: Extracted STIS G750M spectra of the star β Lyrae (exposure rootname o5dh01010) using the standard `calstis` pipeline and various extraction widths as shown in the legend. The unit of the X axis is CCD pixels, and the Y-axis unit is ADU. The scalloped pattern, caused by pixel aliasing, is evident in the uninterpolated spectra.

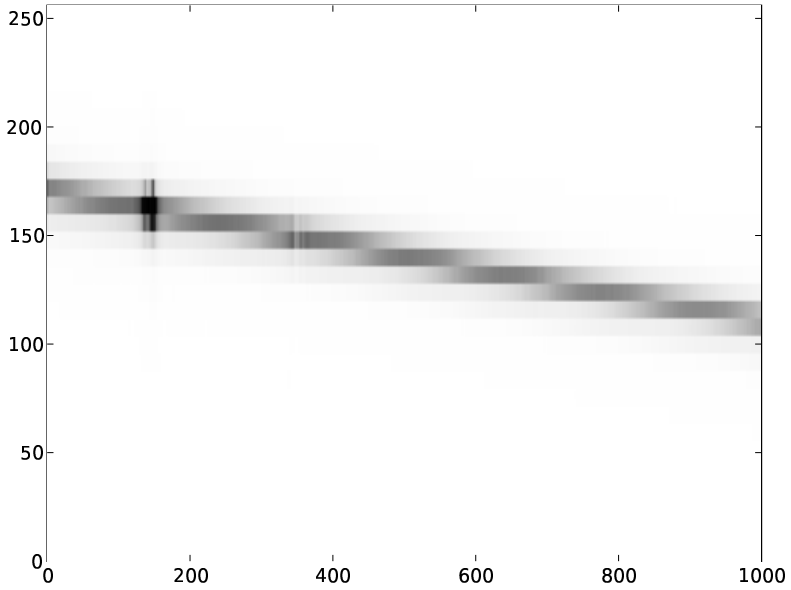

Figure 2: Spectral image of the star whose extracted spectra were shown in Figure 1. The stairstep pattern or aliasing is due to the slightly tilted spectrum and the marginally sampled point spread function. The vertical scale is enlarged to show the aliasing.

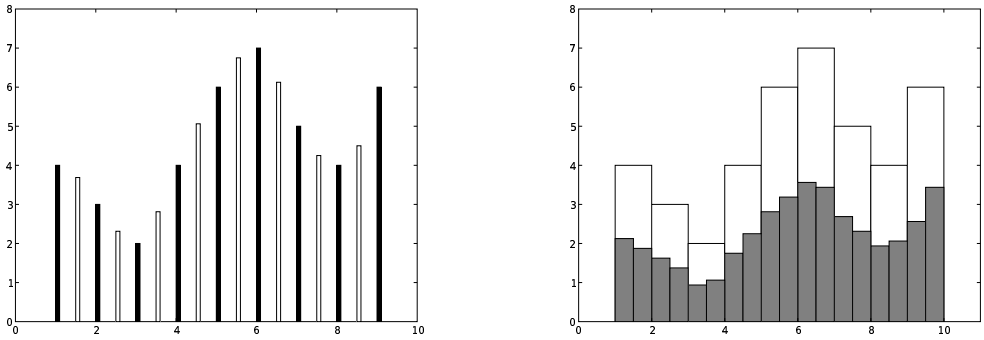

Figure 3: Left panel: Interpolation of data using a 3rd-order polynomial. The original values (filled bars) are located at integral values and the interpolated values (open bars) are at half integral values. Right panel: Average interpolation of data using a 2nd-order polynomial. The original value (white rectangles) span the integers, and the interpolated values (gray rectangles) span the half integers.

difference between adjacent data points is large. Better results can be achieved by using higher (5th or 7th) order polynomials (Figure 3).

Integrating detectors such as CCDs and IR arrays have pixels of finite size. The electrons (or counts) in each pixel are not located at the center of the pixel, but instead are averaged over the area of the pixel. Interpolation — or what we'll call *point-interpolation* — implicitly assumes that the value is located at a point. Instead the underlying function $f(x)$ should be average over the area of the pixel:

$$\lambda_{0,k} = \int_k^{k+1} f(x)dx,$$

where $\lambda_{0,k}$ is the counts in the pixel and k and $k+1$ are the boundaries of the pixel (Figure 3). This is *average-interpolation* and is the most appropriate interpolation for integrating detectors (see e.g. Sweldens 1997, Donoho 1993).

The difference between point-interpolation and average-interpolation may seem minor, but the difference can actually be significant, particularly when accurate fluxes are desired. For astronomical imaging, both interpolation algorithms should give the same total flux for an isolated point source using a sufficiently large extraction region. However, as the extraction region decreases, the measured fluxes will begin to differ. The difference can be $> 50\%$ in some cases with the largest difference occurring where the PSF's gradient is greatest. The reason for this difference is that the distribution of flux within each pixel is different for the two algorithms. This problem becomes more acute when the source is no longer isolated, such as an image of a globular cluster, or the spectrum of an extended source. In this case, the size of the extraction region is limited, possibly by neighboring sources that partially overlap the target. An accurate estimate of the total flux in this case requires an interpolation algorithm that accurately recreates the underlying flux distribution.

3. Wavelets

Wavelets are usually associated with the compression of signals or images by the application of an analysis wavelet (forward transform) of a scaling function. The forward transform uses low and high pass filters to divide the image into smoothed and residual images. The synthesis wavelet is just the inverse transform. It creates a higher resolution image from a coarser image.

Multiresolution is the result of the dilation (or refinement) equation:

$$\phi(t) = \sqrt{2} \sum_{k=0}^{N} c(k)\phi(2t - k),$$

where the scaling function, $\phi(t)$, is a continuous function in t and the coefficients $c(k)$ are discrete in k. This equation relates the scaling function at one scale to the scaling function of the next finer scaler and shows how to calcluate new coefficients that are half way between the old ones. In other words, we can reproduce the scaling function $\phi(t)$ to any resolution by recursion of the dilation equation.

Average-interpolation and (iterative) refinement are two key characteristics of our interpolation algorithm. They ensure that the total counts in a pixel are conserved and the distribution of counts within a subpixel is non-negative. A polynomial is used to approximate the underlying flux distribution. Note that this type of interpolation cannot be done in a single step, i.e. a pixel cannot be subdivided into eight subpixels in one step. Otherwise, there is the possiblity that some subpixels will have negative values, which is not physically possible.

4. Algorithm

The algorithm is a rather simple one-dimensional interpolation problem. The goal is to improve the spatial resolution along the slit (the cross-dispersion direction) in order to reduce the effects of aliasing. We are not interested in improving the spectral resolution, though this is easily achieved by applying this algorithm in the dispersion direction. The partitioning of the counts within each pixel is done by assuming that a N-th order polynomial is a reasonable approximation to the local flux distribution. Average interpolation is used because the CCDs are integrating detectors. Therefore, a polymomial is fit to the cumulative distribution of N-1 pixels, where N is an even integer. The first point of the polynomial is zero, since the integrated area is zero, while the last point is the sum of the N-1 pixels. For each polynomial, only one value, the midpoint of the central pixel, needs to be calculated. Neville's algorithm is used to find this value because it is fast and accurate. The counts in the two subpixels are just the differences between the midpoint value and the points on either side. The algorithm is then repeated for each subpixel until the desired resolution is reached. For STIS three iterations, yielding eight subpixels, usually gives sufficient resolution.

The wavelet algorithm does a very good job of reducing the aliasing, though some may still be present at the few percent level. The aliasing can be reduced further by applying an optional convolution step to the interpolated image. The kernel that is used is an approximation of the instrumental PSF and has the form $(1 + (x/a)^2)^{-2}$, where a is ≈ 1.3 pixels (see also Martin 2004).

5. Results

Figure 4 shows the interpolated image with convolution for three iterations or eight subpixels. The image shows very little of the effects of aliasing. The interpolated spectrum is shown in Figure5 for extraction widths of 1, 2, 4 and 7 pixels. Unlike the uninterpolated spectra, the 1, 2, and 4 pixel width spectra are nearly identical to the 7 pixel width spectrum: only the flux is reduced by a proportional amount.

Figure 6 shows the PSF or cross-dispersion profile for a point source. The black points map the pixel-integrated profile, which is created by shifting the spectral trace so it lies along one row. This causes adjacent pixels in the same row to be shifted slightly with respect to the others, filling in the gaps. The profile is mostly symmetric with wide wings

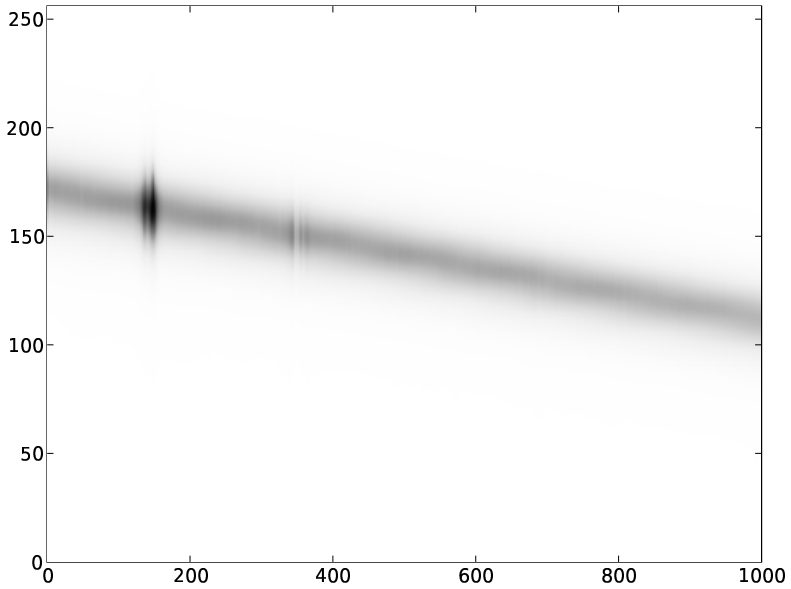

Figure 4: Spectral image of star in Figures 1 and 2 after wavelet interpolation and convolution by the instrumental point spread function.

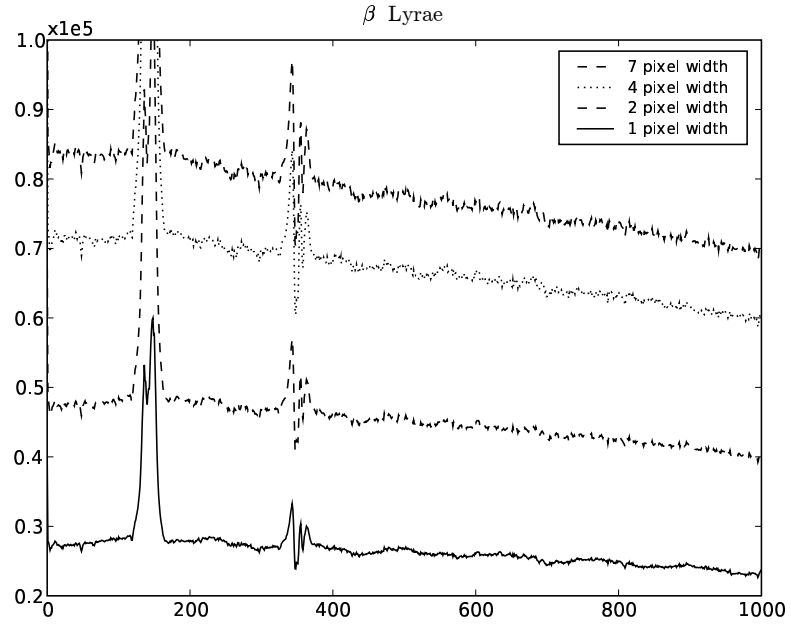

Figure 5: Spectra of star in Figure1 after wavelet interpolation and convolution.

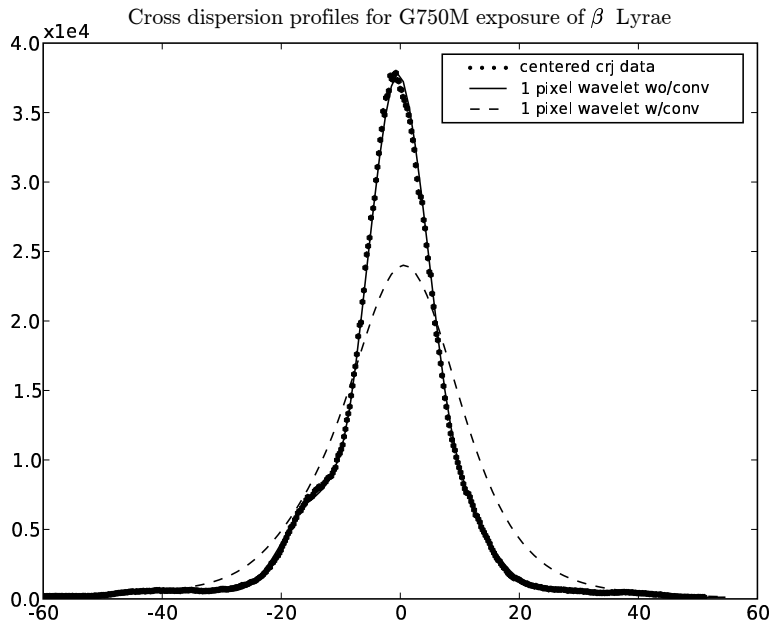

Figure 6: A comparison of the measured pixel-integrated PSF (black points) and the wavelet-interpolated PSF without (solid line) and with (dashed line) convolution.

on each side and a low shoulder on the left side. For comparison, the solid line is the profile along a single column from the interpolated image without convolution. Each point is the sum of eight adjacent subpixels, so that the total area is equal to one pixel as is the case for the black points. A plot of the individual subpixels produces a line that is slightly narrower and more peaked than the pixel integrated profile. The solid line follows the points very closely, except at the peak and near the shoulder. (Note that different columns have different cross-dispersion profiles, depending on the location of the trace: See Figure 9 in Dressel et al. 2006.) The dashed line is the profile from an interpolated image with convolution. This profile is lower and slightly ($\sim 15\%$) broader than the points.

6. Summary

In most situations, interpolation with convolution is the recommended approach. However in situations where the highest spatial resolution is required, such as observations of the binary star WR 140 where the separation of the two stars is about three pixels, interpolation without convolution is preferred, since convolution broadens the image enough to make source confusion a problem. In the case of WR 140, a strong emission line in one stellar spectrum can distort the neighboring spectrum, resulting in incorrect results.

One drawback of this algorithm is that it may become unstable when high order (N$>$ 5) polynomials are used for more than three iterations. In this situation we recommend using a lower order polynomial when doing large number of iterations or beginning with a high order polynomial for the first one or two iterations and then switching to a lower order polynomial for succeeding iterations.

This paper has not discussed what errors are associated with the interpolated values. This aspect of the algorithm will be discussed in a forthcoming paper as will the use of finite impulse response (FIR) wavelet filters on image interpolation or deconvolution.

This algorithm is being added as a post-processing task to the CALSTIS package of STSDAS. PEB wishes to thank Paul Goudfrooij for useful suggestions and the opportunity and time to work on this project.

References

Davidson, K. 2004, Technical memo no. 1, Dept. of Astronomy, Univ. of Minnisota, available through `http://etacar.umn.edu/treasury`.

Dressel, L. L. 2003, in Active Galactic Nuclei: from Central Engine to Host Galaxy, S. Collin, F. Combes, & I. Shlosman (San Francisco: Astronomical Society of the Pacific), 393.

Dressel, L. L., Barrett, P. E., Goudfrooij, P., & Hodge, P. E. 2006, this volume, 267

Donoho, D. 1993, Technical Report, Dept. of Statistics, Stanford Univ.

Kim Quijano, J., et al. 2003, "STIS Instrument Handbook", Version 7.0, (Baltimore: STScI), 240.

Martin, J. 2004, Technical memo no. 2, Dept. of Astronomy, Univ. of Minnisota, available through `http://etacar.umn.edu/treasury`.

Sweldens, W. & Schröeder, P. 1995, SIGGRAPH '95 Course Notes: Wavelets and their Applications in Computer Graphics, 3, 76.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Improving the Rectification of Spectral Images

Linda Dressel

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Paul Barrett

Department of Physics and Astronomy, The Johns Hopkins University, Baltimore, MD 21218

Paul Goudfrooij and Phil Hodge

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. Single rows in STIS rectified spectral images are being analyzed in science programs where the highest obtainable spatial resolution is needed. Here we discuss two problems that have limited the accuracy of the rectified spectral images produced by the STScI pipeline: (1) the changing tilts of the spectral traces across the detector have not been taken into account, and (2) interpolation between rows of spectral images which are undersampled in the cross-dispersion direction has produced substantial artifacts in single rows of the rectified images. To address these problems, post-pipeline software is being developed to produce contemporary traces for science exposures, and the time evolution of traces is being characterized so that it can be incorporated into trace reference files. A technique of iteratively subdividing pixels in the cross-dispersion direction, using area interpolation instead of point interpolation, is being developed for a calibration routine which will produce improved rectified images. Preliminary results of the use of these new techniques will be presented here.

1. Introduction

Accurate spectral traces (which describe the location of spectra projected onto the detector) are especially needed for cases where single rows are to be analyzed in rectified spectral images. Even when the traces are accurate, artifacts are produced by the rectification because the point spread function (PSF) is undersampled in the spatial direction. Combining images dithered by a half pixel along the slit substantially reduces the magnitude of the artifacts. Artifacts can be reduced in spectra obtained from undithered images by using superior methods of interpolation, by smoothing the flux in the cross-dispersion direction, and by using larger extraction boxes. Optimizing the rectification of spectral images involves a trade-off between miminizing the artifacts and preserving spatial resolution.

2. Evolution of Spectral Traces over the Lifetime of STIS

Trace reference files were initially created for STIS using ground-based measurements. Most updates to these files were made in August 1998 using in-flight calibration data obtained by June 1997, a few months after installation. Errors in the traces are often evident in rectified spectral images produced by the pipeline, since individual rows in these images are sensitive to a relative error of a fraction of a pixel in the cross-dispersion direction. It was not known

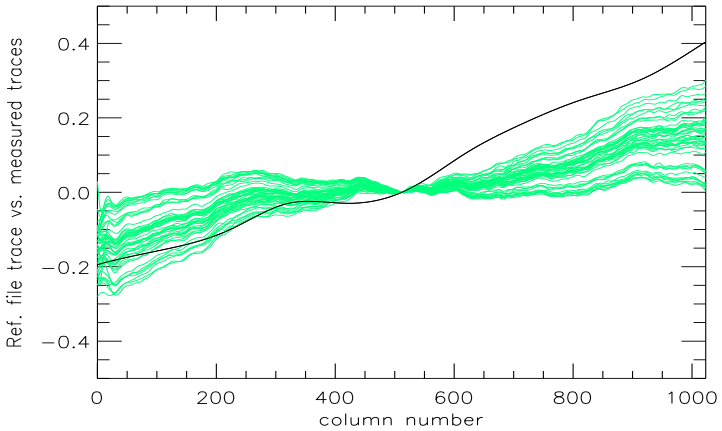

Figure 1: Traces for the G430L grating measured from a star observed at the center of the detector. The smooth dark line is the reference file trace.

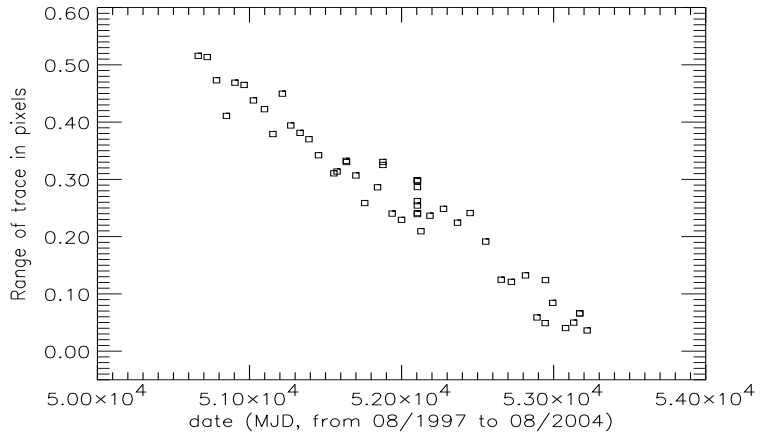

Figure 2: Evolution of the tilt of the trace for the G430L grating. The span of the trace in the cross-dispersion direction is shown as a function of observing date.

at that time whether deviations from the reference file traces were random or systematic with time. As part of the STIS close-out plan (Goudfrooij et al. 2006, this volume), traces are being derived from exposures of stars and other compact sources and examined as a function of observing date for the most commonly used combinations of gratings and central wavelengths.

Figure 1 shows traces measured from spectral images of the standard star AGK+81D266 centered on the detector using the grating G430L with aperture 52×2. The observations were made as part of the sensitivity monitoring calibration programs, and span the operational lifetime of STIS. The reference file trace, also shown, does not quite match any of the observed traces. The full range in y spanned by each trace was measured and plotted as a function of observing date in Figure 2. Clearly, the tilt of the trace depends mostly on time, with much smaller random deviations superimposed on the systematic change. A similar analysis has been done for the gratings G230LB and G750L, with similar results. For all of these gratings, the trace has rotated steadily clockwise with time, changing its range in Y by half a pixel across the 1024 pixels in X.

When the analysis was performed for the grating G750M at the central wavelength setting of 6768 Å, an additional systematic feature was found in the tilt of the traces. Traces of stars observed in calibration and science programs are shown in Figure 3, after

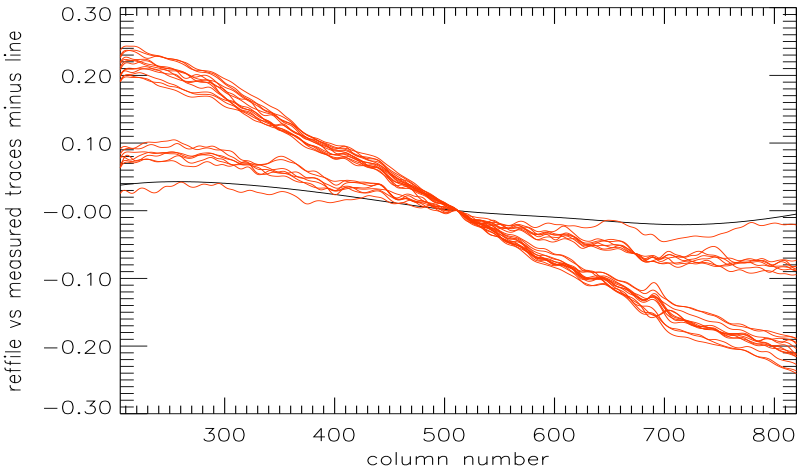

Figure 3: Traces for G750M(6768), with a linear component subtracted, measured from stars observed at the center of the detector. The smooth dark line is the reference file trace.

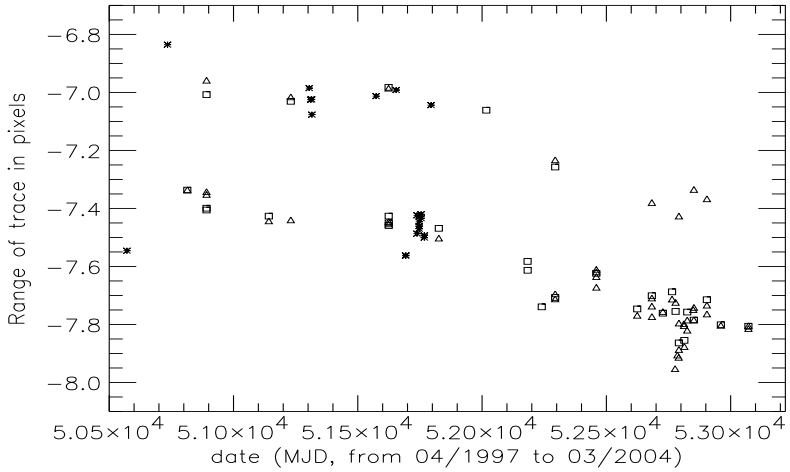

Figure 4: Evolution of the tilt of the trace for G750M(6768). The span of the trace in the cross-disperson direction is shown as a function of observing date. The symbols indicate exposures of stars (*) and long and short exposures (square, triangle) of Eta Carinae.

subtracting the same linear component from each so that the differences in angle can more easily be seen. The traces lie in two bands, but the bands do not represent different time ranges. Most of the observations occurred over a small fraction of the operational lifetime of STIS, so the sample was supplemented with the numerous exposures of Eta Carinae that occurred over a period of many years. The time dependence of the full range in Y spanned by the trace is shown for the entire sample in Figure 4. For this grating, there are two evolutionary tracks separated from each other by about 0.4 pixels in Y range. The same target could have a trace on each track in two visits taken on the same day. We have not determined the cause of this increment. As for the other gratings, each track corresponds to clocekwise rotation of the trace, with a change in Y range of nearly half a pixel over the lifetime of STIS.

For the most frequently used gratings and central wavelengths, we intend to incorporate the time evolution of the tilt of the trace into the SPTRCTAB reference file. In case dual tracks occur for a given grating, the more commonly encountered track will be used. Some science exposures will not be well fit by the reference file trace, either because time evolution

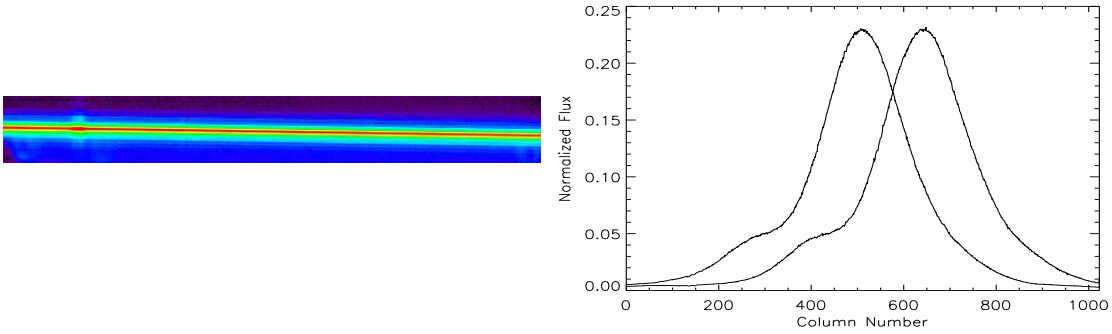

Figure 5: The unrectified spectral image of β Lyrae (left, scaled by 2 in Y) and a plot of two adjacent rows from this image after normalizing out the stellar spectrum (right).

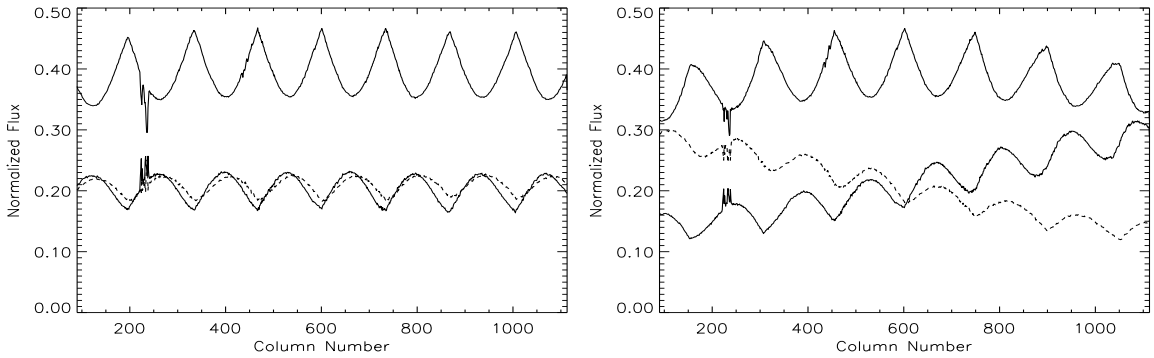

Figure 6: A plot of the peak row and two adjacent rows in the rectified spectral image of β Lyrae after normalizing out the stellar spectrum. The rectification was performed by the stsdas routine x2d using a self-derived trace (left) and the less accurate reference file trace (right).

has not been included for that grating and central wavelength, or because the needed trace lies on the other track of a dual track evolution, or because of random deviations. For those cases, we will supply a PyRAF routine to measure the needed rotation of the trace using selected portions of the science spectrum and to produce the rotated trace.

3. Artifacts in Rectified Spectral Images

Even when traces are well determined, artifacts are produced by the rectification of spectral images. Because the traces are slightly tilted across the detector, interpolation is required in the cross-dispersion direction, and the PSF is undersampled in that direction. The problem is illustrated in Figure 5. The unrectified spectral image of β Lyrae obtained with grating G750M(6768) is shown at the left (exposure rootname o5dh01010). The flux distribution in two adjacent rows of the image is plotted on the right, after dividing each row by the spectrum obtained by summing over all rows. The flux in one row in the rectified spectrum in the wavelength region between the two peaks will be determined by interpolation between these two flux distributions. The interpolated flux at the cross-over point will be lower than the flux at the peaks.

The interpolation performed by `x2d` (a component of CALSTIS), with the spectrum divided out, is shown in Figure 6 for the peak flux row and the two adjacent rows. A self-derived trace was used for the plot on the left; the reference file trace was used for the plot on the right. The broad structure in the latter plot is produced by the error in the tilt of

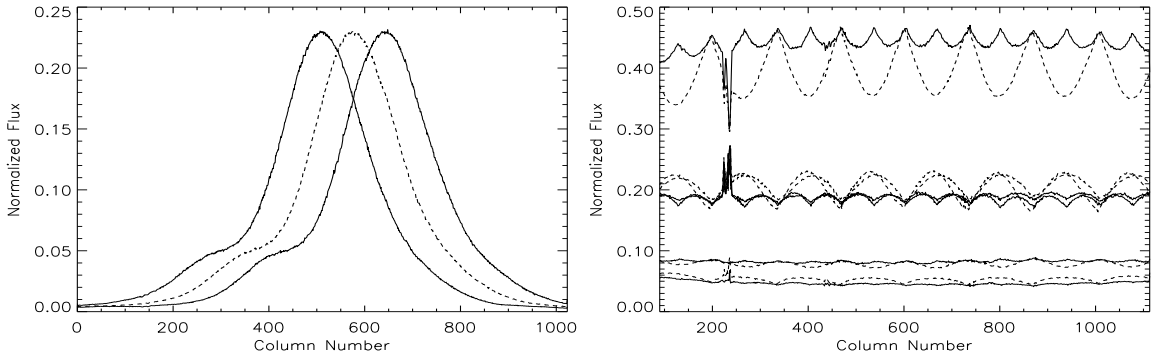

Figure 7: Left: Three adjacent rows from an unrectified synthesized spectral image of β Lyrae created by interleaving the rows in a pair of exposures made with along-the-slit dithering of half a pixel. The stellar spectrum has been normalized out. Right: A plot of the peak row in the rectified image and the rows offset by one and two CCD pixels. Rows from the synthesized image (solid lines) are compared to rows from a single image (dashed lines).

the applied trace, whose total range in Y is off by 0.7 pixels. The `x2d` parameter SHIFTA2 was adjusted to center the output spectrum on a row. When the trace is accurate, this preserves the peaks seen in each row in the pre-rectified image, since they are shifted by integer pixels. Between these peaks, the interpolation places too little flux in the peak row and too much flux in the adjacent rows. There are many cycles of fluctuations across the spectrum for gratings with the most tilted traces: G230L (18), G140M (10), G230M (16), G430M (8), and, shown here, G750M (7). A 3-row spectral extraction greatly reduces the magnitude of the artifacts since they are out of phase in adjacent rows.

Rectification artifacts can be significantly reduced observationally by using a half-pixel dither along the slit to provide better sampling in the spatial dimension. To achieve the improvement, the rows of the two spectral images are interwoven before rectification. Figure 7 (left) shows how a row from the dithered image (rootname o5dh01020) has been inserted between the two rows from the original image plotted in Figure 5, with the stellar spectrum again normalized out. The rectification is performed using a trace derived from the interwoven image. Figure 7 (right) shows the resulting peak row and the rows at offsets of one and two true pixels (solid lines). The analogous rows from a single image are also plotted (dashed lines). Fluxes are preserved by integer shifts along the slit in twice as many locations in the interleaved image, and the interpolation losses from the peak row at the intermediate locations are much less. The amplitude of the fluctuations is reduced by a factor of 3.4 in the peak row, and by a factor of 2.8 in the adjacent rows.

For an unresolved source, the artifacts of rectification can easily be seen by normalizing out the spectrum, as shown in the example above. For spatially complex sources, real spatial variations in the spectrum are mixed in with the artifacts in the rectified image, and cannot be simply normalized away. In this case, the rectification product of interwoven dithered images can be compared to that of a single image to illustrate the effect of the undersampling artifacts on the spectrum. Here we consider the spectral imaging of NGC 3998, a LINER galaxy with compact and extended emission lines, a compact power-law continuum, and compact and extended stellar emission. As for Beta Lyrae, a half-pixel dither along the slit was performed between exposures. Figure 8 (left) shows three adjacent rows from the interwoven spectral image prior to recification. The intermediate row taken from one exposure (dashed line) clearly differs from an interpolation of the nearest rows taken from the other exposure (solid lines). A row in the rectified image is shown for a single input image (dashed line) and the combined interwoven images (solid line) in Figure

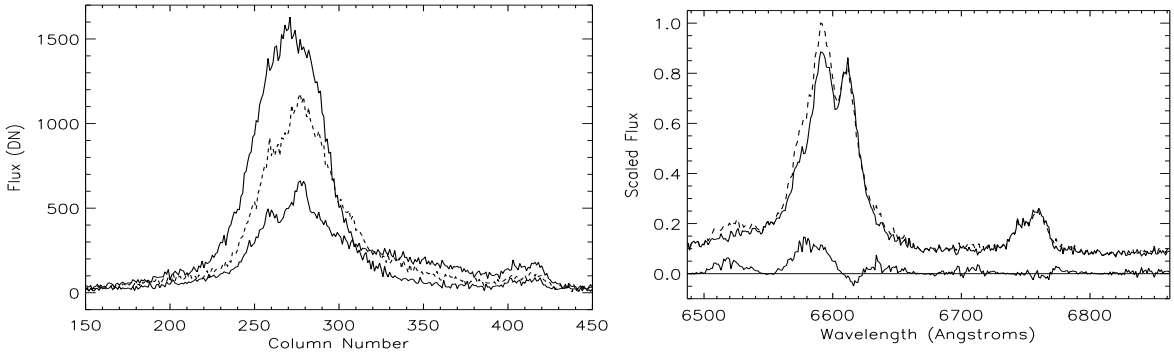

Figure 8: Left: Three adjacent rows from an unrectified synthesized spectral image of the LINER nucleus of NGC 3998 created by interleaving the rows in a pair of exposures made with along-the-slit dithering of half a pixel. Right: A plot of a row in the rectified synthesized image (solid line) and the analogous row in a rectified single image (dashed line), scaled to peak 1.0. The difference between these two spectra (solid line) shows the undulations produced by compact components in the line and continuum emission, mostly due to spatial undersampling in the single image.

8 (right). The difference, due mainly to the rectification artifacts of the single image, is also shown (solid line). This difference clearly affects the measurement of the parameters of the emission lines. The rotation curve derived from the apparent emission-line velocities in the single image was asymmetric and not centered on the photometric center, whereas the rotation curve derived from the interwoven dithered images *was* symmetric and centered on the photometric center, in good agreement with the standard model of a rotating gaseous disk centered on a supermassive black hole (Dressel 2003).

4. The Rectification Process Using Wavelet Interpolation and Convolution

Many STIS spectral imaging programs did not use along-the-slit dithering to improve the sampling of the PSF in the spatial direction. For these programs, improvements in rectification depend on improvements in the data processing techniques. Here we present wavelet interpolation (Barrett & Dressel 2006, this volume) as a superior method, to the bilinear interpolation used in the STIS pipeline when used near the center of compact sources. An alternative method is presented by Davidson (2006, this volume).

The effects of undersampling can be seen clearly in the upper panel of Figure 9, which shows flux profiles of β Lyrae along the slit for a column where the observed trace is centered on a row and a column where it straddles two rows, prior to rectification. The two columns have a small separation in wavelength, so the PSF striking the detector is virtually identical at these columns. (Charge diffusion on the CCD contributes modestly to the difference in the observed flux profiles. See Dressel 2006 for subpixel modelling and measurement of the PSF.) The lower panel of Figure 9 shows how the flux within each pixel is redistributed into 8 subpixels using the wavelet interpolation algorithm developed by Barrett (Barrett & Dressel 2006). Also shown is the result of convolving each subpixeled profile by the PSF kernel defined in Barrett & Dressel (2006), adopted from the approximation developed by Davidson (2004). Rectification is performed by regrouping the subpixels into pixels that are centered on the trace. Wavelet interpolation improves the agreement between the subpixeled profiles, but it is the convolution that makes the profiles nearly identical. Indeed, if each pixel is simply divided into 8 subpixels with the *same* flux, the convolved profiles are nearly identical to those in Figure 9. The periodic fluctuations in the rows of the rectified spectrum

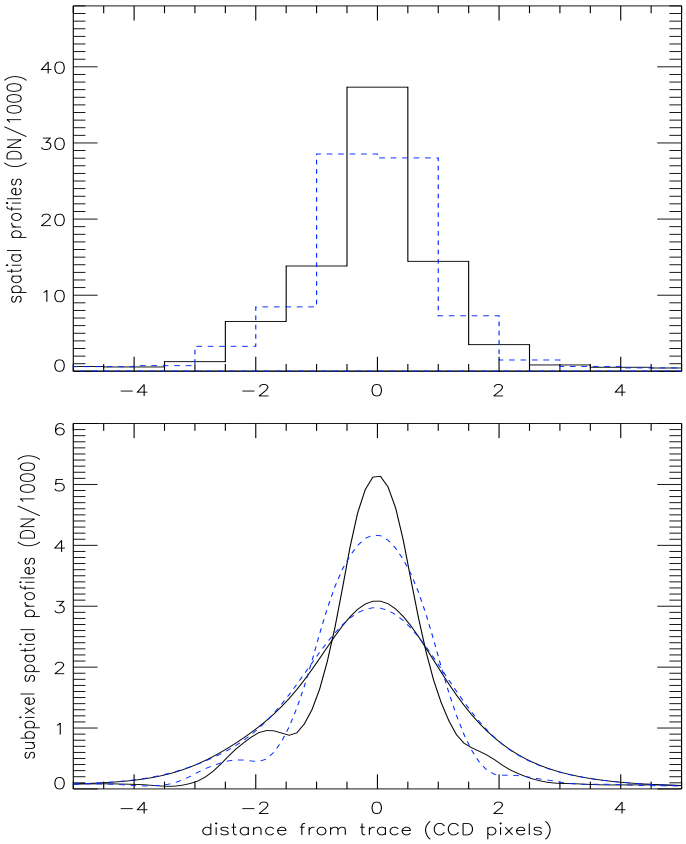

Figure 9: *Upper panel*: Along-the-slit flux profiles of β Lyrae in two columns of the unrectified spectral image, where the trace is centered on a row (solid line) and where it straddles two rows (dashed line). *Lower panel*: Redistribution of the flux in each column using recursive wavelet interpolation (the two narrower profiles) followed by convolution with a PSF kernel (the two very similar broad profiles).

disappear as the spatial profiles are convolved to nearly the same shape, but this is clearly achieved at the expense of spatial resolution.

Spectral extractions produced using convolution by the PSF kernel are shown in Barrett & Dressel (2006) for β Lyrae. Here we show spectral extractions produced without convolution. The result of wavelet interpolation, with the spectrum divided out, is shown in Figure 10 for the peak flux row and the surrounding rows (solid lines). The analogous spectra are shown for the interpolation performed by `x2d` (dashed lines). Wavelet interpolation reduces the amplitude of the fluctuations in the peak row by a factor of 1.6. Wavelet interpolation and `x2d` give similar results for the rows adjacent to the peak row. In the next row further out, the fluctuations are larger in the wavelet-intepolated spectra than in the `x2d` spectra.

We return now to the spectral image of the LINER galaxy discussed above, rectifying a single exposure using `x2d`, wavelet interpolation, and wavelet interpolation with convolution. The results will be compared to the `x2d` rectification of the interwoven dithered images, which achieves lower-amplitude fluctuations in the spectrum through improved spatial sampling as demonstrated above. The peak row in the rectified image is shown in Figure 11 (upper panel) for the interwoven images, the single `x2d` image, and the wavelet-interpolated image without convolution by the PSF kernel. The wavelet-interpolated spectrum is intermediate between the `x2d` products of the single and interwoven images. The

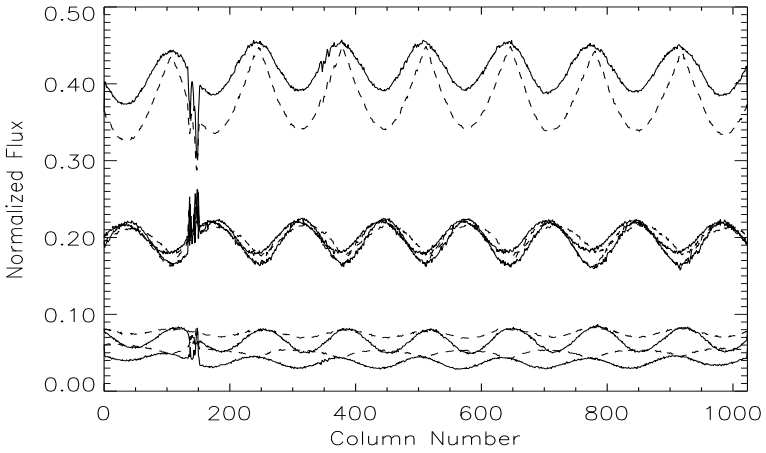

Figure 10: The peak 5 rows in the spectral image of β Lyrae (with the total spectrum divided out) produced by wavelet interpolation (solid lines) and by `x2d` (dashed lines).

difference between the wavelet-interpolated spectrum and the interwoven `x2d` spectrum is also plotted. Most of this difference is due to the greater magnitude of the artifacts in the wavelet-interpolated spectrum. One row below the peak row (Figure 11, central panel), the wavelet-interpolated spectrum is similar to the single-image `x2d` spectrum, as expected from Figure 10, and significantly different from the interwoven-image `x2d` spectrum. Two rows below the peak row (Figure 11, lower panel), the wavelet-interpolated spectrum alternately matches the single-image `x2d` spectrum and dips below it (again consistent with Figure 10) and differs significantly from the interwoven-image `x2d` spectrum. Figure 12 is analogous to Figure 11, except that the wavelet-interpolated image has been convolved with the PSF kernel in the cross-dispersion direction, resulting in substantial redistribution of flux along the slit. In all cases, one can see that fits of models to the emission lines (Hα with a blue wing and [NII]) will be affected by the method used to produce the rectified spectral image.

5. Conclusion

Spectral traces show time evolution that will be incorporated into reference files for some commonly used gratings and central wavelengths. We are developing post-pipeline software to generate traces for less commonly used modes and modes with additional complications. Single rows in rectified images must be used carefully, with a full understanding of the artifacts. Along-the-slit dithering is the most effective way to enhance spatial resolution in spectral images. Wavelet interpolation and convolution in the spatial dimension can be helpful in reducing the artifacts in rectified spectral images. The impact of spectral artifacts must be weighed against the need for spatial resolution in deciding how to apply them.

References

Barrett, P., & Dressel, L., 2006, this volume, 260

Davidson, K. 2004, Technical Memo #1, http://etacar.umn.edu/treasury/techmemos

Davidson, K. 2006, this volume, 247

Dressel, L. 2003, in Active Galactic Nuclei: from Central Engine to Host Galaxy, S. Collin, F. Combes, & I. Shlosman (San Francisco: Astronomical Society of the Pacific), 393

Dressel, L. 2006, this volume, 267

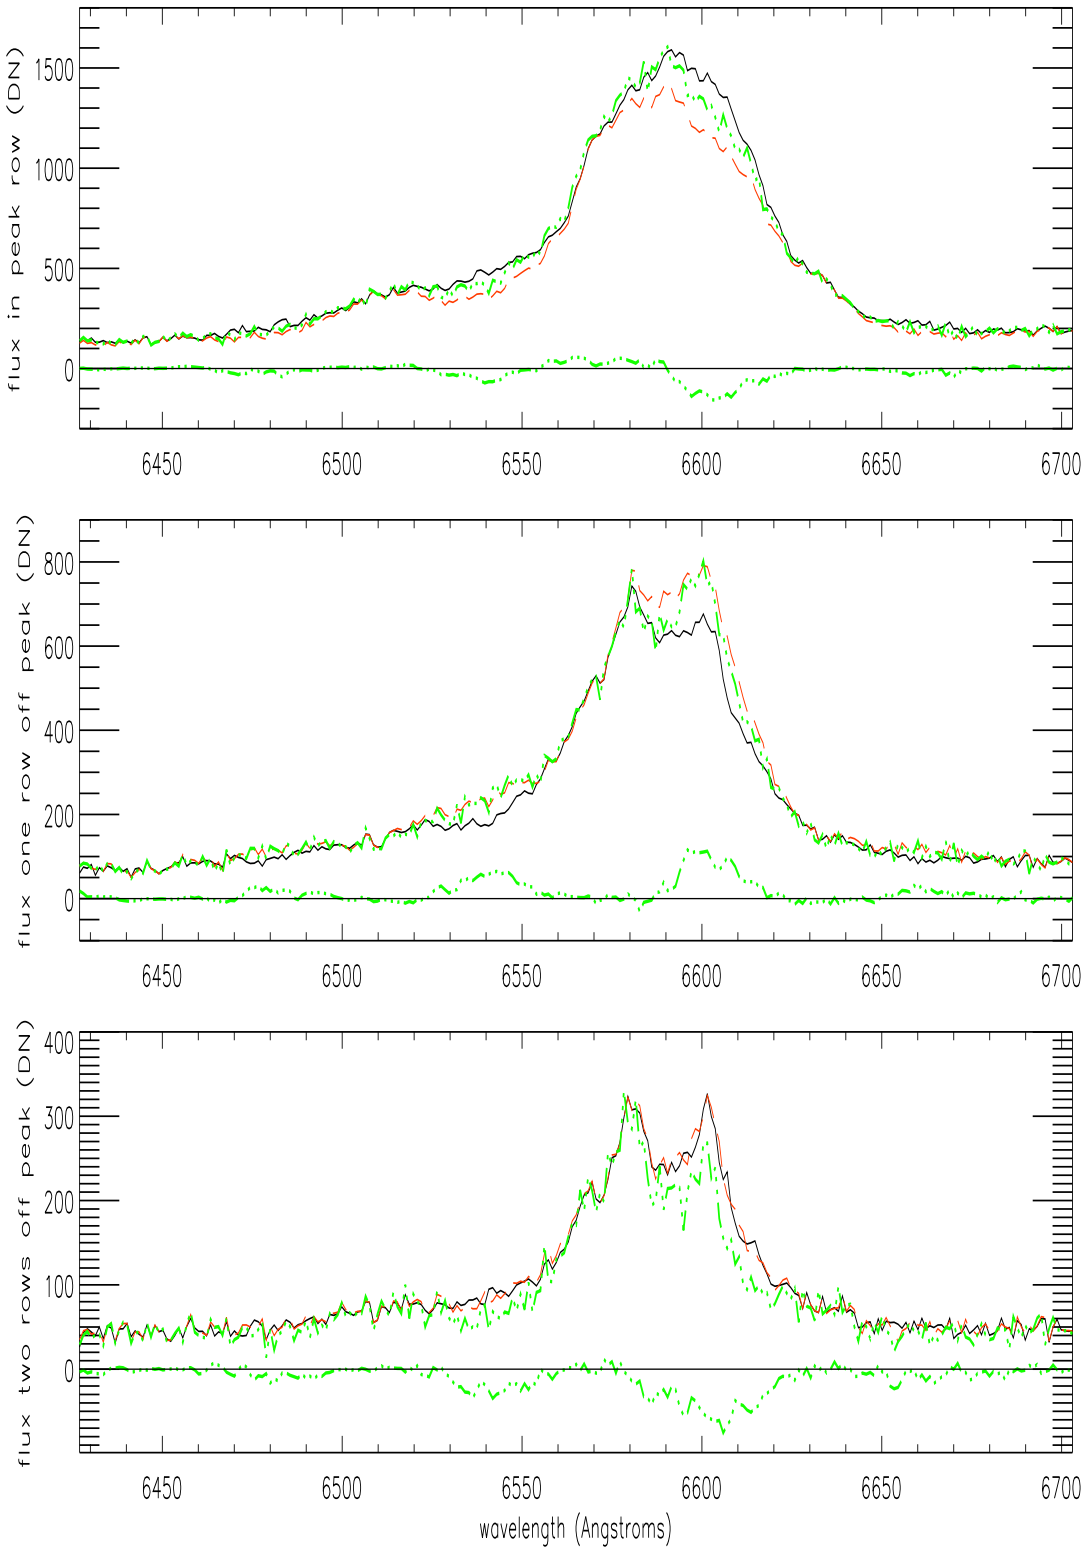

Figure 11: Rows from rectified spectral images made by applying `x2d` to a single image (dashed line), `x2d` to interwoven dithered images (solid line), and wavelet interpolation to a single image (dot-dashed line); the difference between the last two is also shown (dot-dashed line). Upper panel: peak row; middle panel: one row below peak row; lower panel: two rows below peak row.

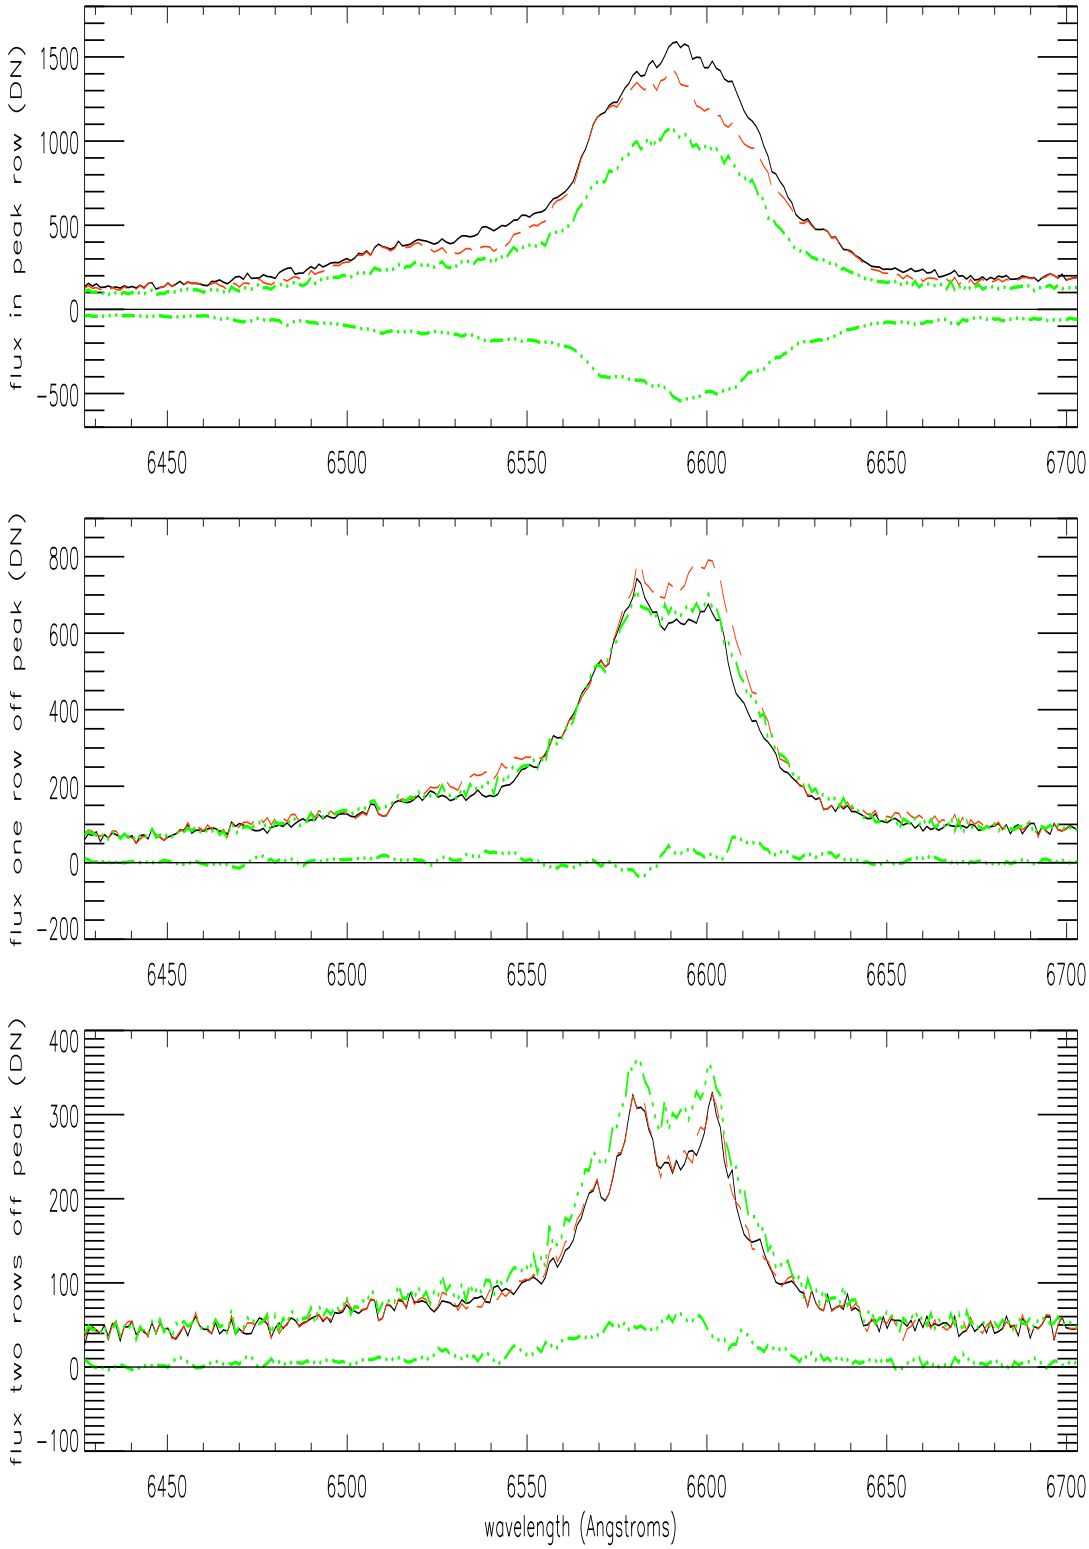

Figure 12: Same as Figure 11, except that the wavelet-interpolated image has been smoothed in the cross-dispersion direction by the PSF kernel.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Spectroscopic Point Spread Functions for Centered and Offset Targets

Linda Dressel

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD, 21218

Abstract. Finely sampled observational spectroscopic point spread functions (PSFs) have been generated from exposures of stars. Lyot stops near the pupil plane of the STIS G430L and G750L gratings block some scattered and diffracted light at the expense of reducing the throughput and broadening the PSF. For G750L, perpendicular-to-slit stepping patterns were performed across a star with the 52x0.1, 52x0.2, 52x0.1E1, 52x0.2E1 apertures to measure the spectroscopic PSF for an out-of-slit target, relevant to spectoscopic mapping observations. The observed spectroscopic PSFs have been compared to the predictions of Tiny Tim models, which were developed for direct imaging. Models can easily be generated using no Lyot stop or using the Lyot stop parameters for the CCD mirror, similar to those for the G450L and G750L gratings. The on-target PSFs were generally well fit except for excess observed scattered light at very low levels. The observed offset-slit G750L fluxes were less than predicted, even after adjusting for the difference in size between the PSF sampled in the slit plane and the PSF modelled on the detector. Observed and modelled PSFs for G750M (with no Lyot stop) and G750L (with a Lyot stop) at the same wavelength are compared.

1. Modelling and Measuring the Spectroscopic PSF

To model a spectroscopic PSF, one begins by producing a finely sampled monochromatic imaging PSF using Tiny Tim software (Krist and Hook, http://www.stsci.edu/software/tinytim). The imaging PSF appropriate to the grating must be used, generated with a Lyot stop (for G430L or G750L; Heap et al., 2000) or without a Lyot stop (for the other first order gratings). To generate the spectroscopic PSF, the aperture is placed on the imaging PSF and the flux within the aperture is summed along the dispersion direction. The column of summed fluxes is blocked into pixels in the cross-dispersion direction. For CCD modelling, the column is convolved with a one dimensional kernel to simulate charge diffusion on the CCD. The kernel is obtained by collapsing the Tiny Tim kernel to a single dimension, as appropriate for a locally flat or normalized spectrum. A PSF continuously sampled along the slit is shown in Figure 1, before blocking the subpixels into pixels, after the blocking, and after applying charge diffusion. For comparison to an observed column of flux in a spectral image, one must choose points on the charge-diffused profile at intervals of 0.05 arcsec (one CCD pixel).

Two techniques can be used to produce finely sampled observational spectroscopic PSFs, depending on the tilt of the spectral trace:

1. For gratings with a spectral trace that drops by several rows as it crosses the detector, a single spectral image is sufficient to produce a finely sampled PSF. The fractional pixel drops in the trace from one column to the next in an flt or crj image can be treated as a series of small dithers. The observed PSF is produced by normalizing out the stellar spectrum and centering the flux profile in each column on the trace. For the M gratings, a band of many columns can be used to sample the PSF because it does not measurably

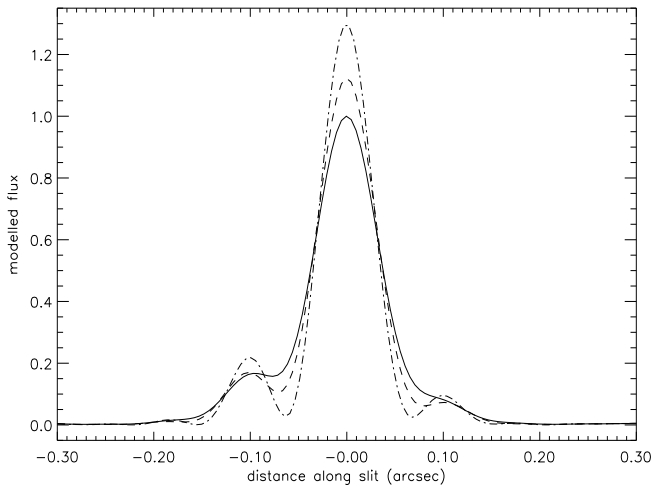

Figure 1: The spectroscopic PSF at 6600 Å for the 52x0.1 slit before blocking the subpixels into pixels (dot-dashed line), after the blocking (dashed line), and after applying charge diffusion (solid line).

change across the short span of wavelengths. A spectral image of the star BD+75D325 taken with the 52X2 aperture was used to produce the G750M PSF at 6600 Å.

2. For gratings with a spectral trace that is nearly flat, the slight randomness in the placement of the spectrum on the detector can be treated as a dither for a sample of many exposures. The flux profile from a single column in each image, centered according to its placement on the detector, can be used to measure the PSF at a given wavelength. Spectral images of the star AGK+81D266 from the STIS sensitivity monitor programs were used to produce the 52X2 G750L PSF at 6600 Å.

A comparison of the observed 52X2 6600 Å PSFs for G750M (without Lyot stop) and G750L (with Lyot stop) is shown in Figure 2. The profiles expected from Tiny Tim modelling are also shown. The effect of the Lyot stop on the PSF in the cross-dispersion direction is clearly seen in the profiles, and the modelling shows good consistency with the data for both gratings. The most conspicuous difference between the G750L and G750M profiles is the broader "shoulders" of the G750L PSF. As seen in Figure 1, coarse sampling by pixels and charge diffusion on the detector can obscure the source of the difference between the spectroscopic PSFs. The source of the difference in the modelling is an overall expansion of the imaging PSF by the Lyot stop. When the subsampled imaging PSF without Lyot stop is expanded by 10% with flux conservation, its radial profile closely matches that of the subsampled imaging PSF with Lyot stop.

2. Sampling the PSF with Narrow Slits

The 52X0.1 and 52X0.2 slits were frequently used in perpendicular-to-slit stepping patterns for the purpose of mapping the spatial distribution and kinematics of extended sources. To assess the accuracy of the Tiny Tim modelling used in the analysis of such data, observations of HD73471 with grating G750L were obtained in STIS calibration program 9610. The 52X0.1 slit was stepped in a perpendicular-to-slit pattern of five 0.1 arcsec steps centered on the star. Similary, the 52X0.2 slit was stepped in a perpendicular-to-slit pattern of three 0.2 arcsec steps centered on the star. Both patterns were repeated at the E1 aperture positions, which place the target near the readout end of the CCD detector to reduce CTI losses. For each slit position, four dither steps were performed along the slit at 3.5 pixel

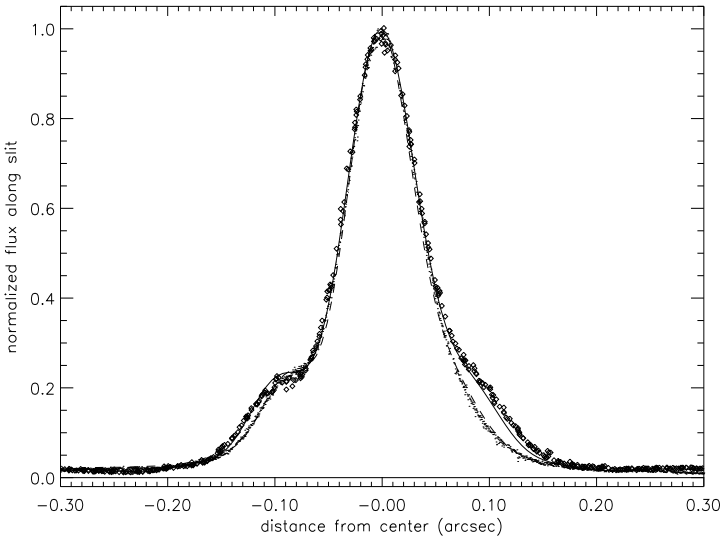

Figure 2: 6600 Å PSFs for 52X2 G750L (data: diamonds, model: solid line) and G750M (data: dots, model: dashed line).

intervals to give half-pixel sampling and redundant pairs of observations. A single column at 6600 Å in each spectral (flt) image was used to form the observational PSF. The slit placements are shown in Figure 3, where the star is represented by the Tiny Tim imaging PSF (with Lyot stop) subsampled to 0.1 pixel (0.005 arcsec). The trefoil structure in the first Airy ring is caused by the three support pads on the HST primary mirror.

Models of the spectroscopic PSF were produced as described in Section 1, with one modification. Tiny Tim generates models of the PSF as it appears in the plane of the detector. The expansion of the PSF by the Lyot stop must be taken into account so that sampling of the PSF in the aperture plane can be modelled appropriately for narrow slits. The slit has been broadened by 10% when applied to the Lyot-stopped imaging PSF, to compensate for the 10% narrower PSF in the aperture plane. (This broadening has been applied to the slits shown in Figure 3.) The detailed features of the Lyot-stopped PSF are thus preserved, and the extent of the PSF in the cross-dispersion direction at the detector is maintained, but the sampling by the slit is improved.

For each slit, the subsampled (0.1 pix) PSF was summed over grids representing slitwidth times pixel height. Since there is no constraint on the y positioning of the spectrum within a pixel, the summation was performed 10 times, stepped by one subpixel in the y dimension each time, to give 10 sets of spectroscopic PSFs to compare to the data. For each slit position and y-centering, the column of numbers representing the PSF along the slit was then convolved with the charge diffusion kernel. The modelled profiles with the y-centering that best matched the data in the central slit were selected. The data were normalized to have the same flux summed over the central 0.25 arcsec in the central slit as the model.

Figure 4 shows the observed flux along the central slit and the PSF model prediction for the 52X0.1 slit. The left column of plots shows the results for the 52X0.1 observations (centered on the detector), and the right column shows those for the 52X0.1E1 observations (high on the detector). The three rows of plots show different ranges in intensity and distance along the slit. The fluxes are somewhat lower than predicted in the brighter segment of the Airy ring. They are greater than predicted in the faint wings of the PSF, due to the halo of scattered light not included in the Tiny Tim modelling.

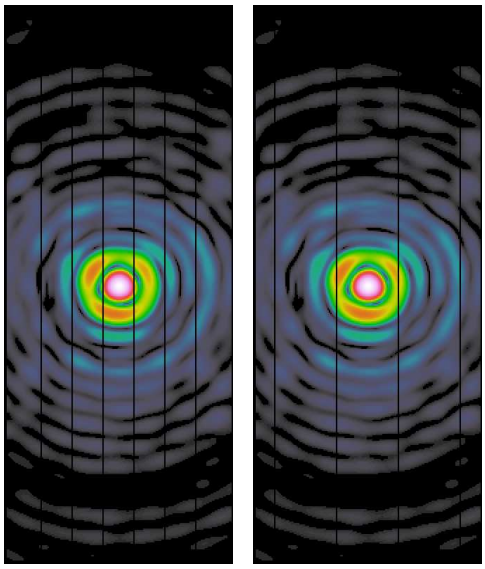

Figure 3: Model 6600 Å PSF subsampled to dimensions of 0.1 pixel, with overlays of the 5 slit placements of the 52x0.1 slit (left) and the 3 slit placements of the 52x0.2 slit (right). The slits have been broadened by 10% to compensate for the broadening of the PSF caused by the Lyot stop.

Figure 5 shows the central slit flux profile for 52X0.1 again, modelled and observed, in the central panel. The upper and lower panels show the flux profiles for the intermediate and outer slit positions in the 5-step observing pattern. The modelled profiles in the offset slits generally over-predict the flux. A similar result was found for the offset slits in the 3-step 52X0.2 observing pattern. The model predictions are sensitive to the position and structure of the first Airy ring, as can be seen in Figure 3.

3. Summation of Slits in Patterns

An assumption of the modelling is that the summation of the observed fluxes in the five contiguous positionings of the 52X0.1 slit should equal the flux observed in the 52X0.5 slit. The summation of the observed fluxes in the three contiguous positionings of the 52X0.2 slit should be similar to the flux in the 52X0.5 slit, since the flux near the outer edges of the 52X0.5 slit is small. The profiles of the 6600 Å flux along the slit are shown in Figure 6 for the summed 52X0.1E1 observations, the summed 52X0.2E1 observations, and a single 52X0.5E1 observation of the same star made one year earlier. (Fluxes from dither positions separated by half a pixel along the slit are shown for the summed observations; the 52X0.5E1 observation was not dithered.) The summed 52X0.2E1 fluxes nearly equal the 52X0.5E1 fluxes, as expected, but the summed 52X0.1E1 fluxes are noticeably lower.

Several potential causes of the low summed 52X0.1 fluxes, and of the discrepancies between modelled and observed fluxes in the offset slits, can apparently be ruled out. The focus was known and taken into account in the modelling, and changes in focus due to breathing were insignificant during these observations. ACQ/PEAKs ensured accurate pointing, and errors in the small angle maneuvers executed in the stepping patterns are small (Kim Quijano et al. 2003). The slit sizes have been well measured (Bohlin and Hartig 1998). Temperature-dependent sensitivity changes for this grating are neglibible (Stys et al. 2004). The size of the PSF in the aperture plane cannot be much less than assumed. A remaining

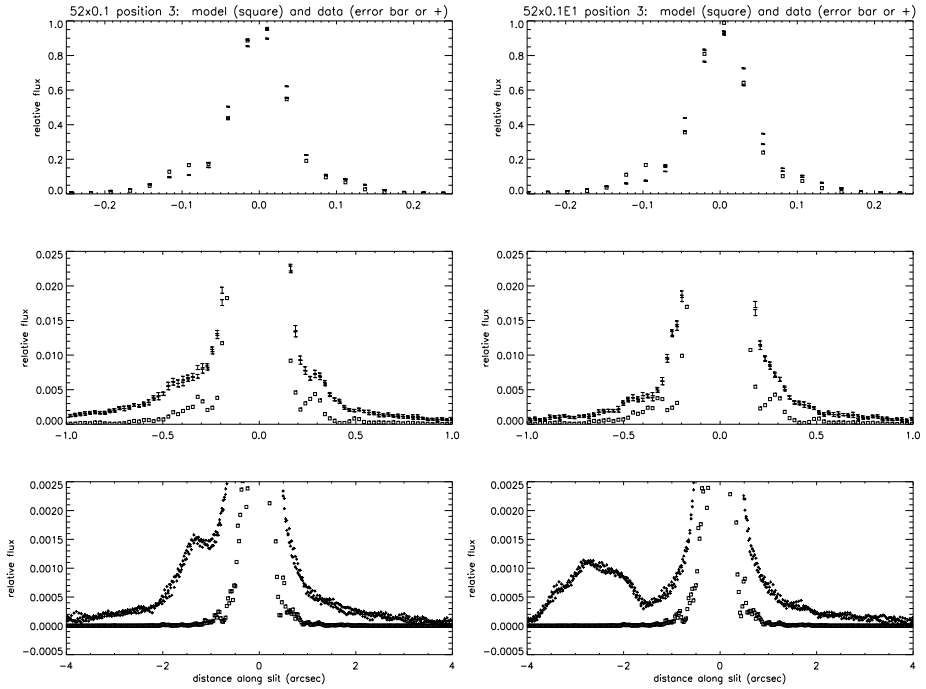

Figure 4: Observed flux along the slit (error bars or +) and PSF model prediction (squares) for the 52x0.1 aperture (left) and the 52x0.1E1 aperture (right) centered on the star. Three ranges in distance along the slit and in flux are shown.

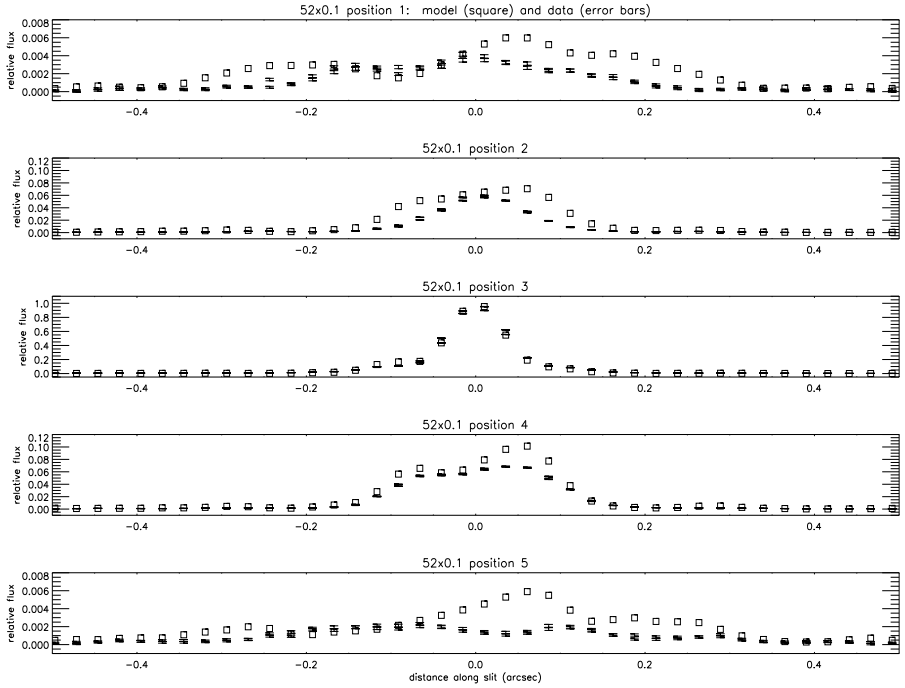

Figure 5: Observed profiles (error bars) and modelled profiles (squares) of flux along the slit for each of the 5 placements of the 52x0.1 aperture.

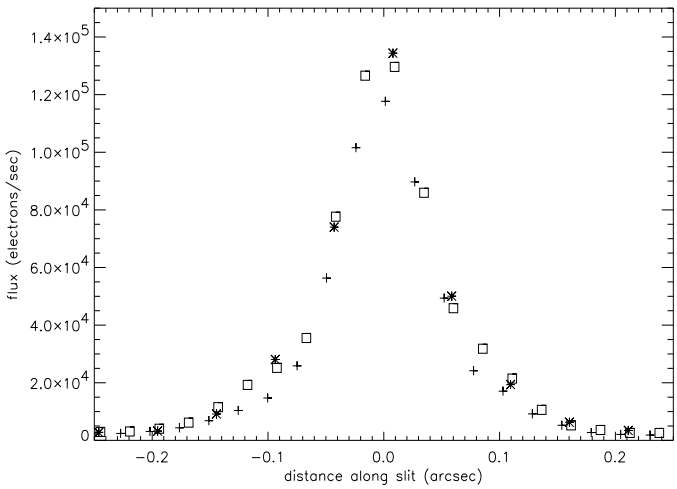

Figure 6: Observed spectroscopic PSFs: summation of fluxes in 5 contiguous positionings of aperture 52x0.1E1 (+), summation of fluxes in 3 contiguous positionings of aperture 52x0.2E1 (square), and flux in a single position of aperture 52x0.5E1 (*).

source of error may be the treatment of the effect of the Lyot stop on the slitted PSF, using only a scaling factor for the slit width.

4. Conclusions

I have compared Tiny Tim model predictions and G750L flux profiles at 6600 Å for a star centered in a slit (52X0.1, 52X0.1E1, 52X0.2, or 52X0.2E1) and moved out of the slit by one or two slit widths. The model under-predicts the flux in the centered slit at distances greater than 0.2 arcsec from the target, where it does not fully account for scattering and diffraction. It over-predicts the flux in the offset positions, possibly due to the simplified treatment of the effect of the Lyot stop on slitted light. The G750M grating does not have a Lyot stop, and therefore has a narrower PSF than the G750L grating. The observed PSFs for these gratings are differently shaped because of the combined effects of broadening by the Lyot stop, undersampling by pixels, and charge diffusion on the CCD. For a very broad slit (52x2) centered on the star, Tiny Tim modelling reproduces the observed spectroscopic PSF well for both gratings. A more detailed presentation of this analysis will be given in a STIS ISR.

Acknowledgments. I thank Chuck Bowers, Paul Goudfrooij, Ted Gull, George Hartig, Don Lindler, and Charles Proffitt for useful discussions and information.

References

Bohlin, R., & Hartig, G. 1998, STIS *Instrument Science Report* 1998-20 (Baltimore: STScI)

Heap, S.R., Lindler, D.J., Lanz, T.M., Cornett, R.H., Hubeny, I., Maran, S.P., & Woodgate, B. 2000, ApJ, 539, 435

Kim Quijano, J., et al. 2004, *STIS Instrument Handbook* (Baltimore: STScI)

Stys, D.J., Bohlin, R.C., & Goudfrooij, P. 2004, STIS *Instrument Science Report* 2004-04 (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Improving the Spectral Resolution of HST Spectra via LSF Replacement

Steven V. Penton
Center for Astrophysics and Space Astronomy, University of Colorado, Boulder, CO 80309

Abstract. Large aperture spectra provide greater signal, but small aperture spectra have higher resolution (smaller line-spread function (LSF) wings). The simple method presented here can convert large aperture spectra to the resolution of small aperture spectra by removing the wings of the LSF without significant noise being added to the continua. In addition, this method/algorithm allows spectra taken through different apertures to be placed upon a common baseline for increased signal-to-noise (S/N). In marginal S/N cases, such as QSO absorption studies, this method is a particularly promising archive mining tool.

1. Introduction

The method is inspired by the 2D 'CLEAN' method used in radio astronomy (Hőgbom, 1974), but adapted to 1D spectra. At the location of the highest flux point in the spectrum, a wavelength-dependent line-spread function (LSF) of photons is subtracted and its location stored. The LSF of photons subtracted is a small fraction ($< 10\%$) of the number of photons at this spectral location. Simultaneously, a new spectrum is created by adding a replacement LSF of photons at the original spectral location in the new spectrum. The replacement LSF could be a Gaussian of specified width, a delta-function, the LSF of another grating/aperture combination, or any other LSF (such as a 'wingless' approximation of the original or alternate LSF). Successive passes through the remaining spectrum result in a residual spectrum that, after many such passes, begins to resemble a featureless continuum. Iterations are truncated when the maximum remaining flux is a predetermined fraction of the original error vector, or, equivalently, when a certain size feature is no longer present in the spectrum. After this process has finished, a new spectrum is created by adding the remaining residual spectrum to the "cleaned" version.

2. Line-Spread Functions (LSFs)

The LSF describes the instrumental spectral distribution of an incident monochromatic emission source (delta-function) and is a function of aperture, grating, and wavelength. Sample HST/STIS first-order MAMA LSFs are shown in Figure 1 (taken from the STIS Instrument Handbook, Quijano et al. 2004). With increasing aperture size, and photon energy, the percentage of flux in the wings of the LSF increases significantly. HST spectral observers are forced to consider the trade-offs between apertures, often sacrificing flux (signal-to-noise) for resolution. For example, even though the reported literal full width half maximums (FWHM) are nearly identical, at 1200Å G140M observers were forced to select from apertures ranging from 42 to 91% transmissions and undocumented spectral resolution changes of greater than 20%.

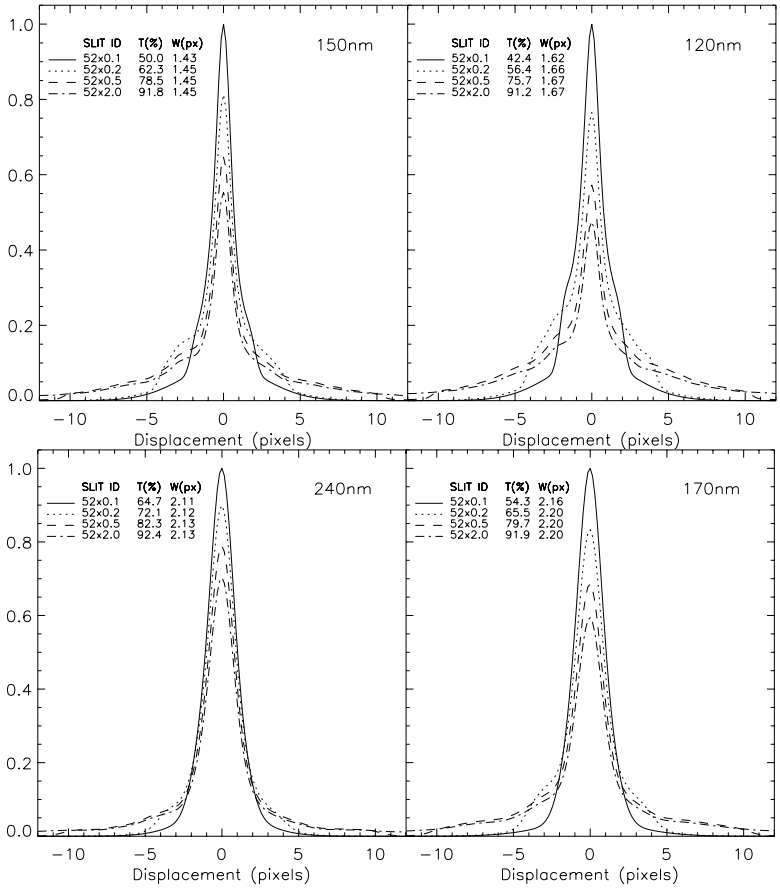

Figure 1: STIS first-order MAMA LSFs for the G140L and G140M (top) and G230L and G230M (bottom). The legend in the upper left gives the aperture, SLIT ID, the percentage of the incident flux transmitted by the aperture, T(%), and the FWHM of the central core in pixels, W(px). The wavelength is indicated in the upper right corner. The spectral method presented here is designed to increase spectral resolution by removing the significant, broad non-Gaussian wings.

3. Resolution Enhancement

Spectral resolution (R=$\lambda/\Delta\lambda$) is often defined by the literal FWHM of the central Gaussian core of a LSF and is only achievable for delta-function absorption or emission lines. To quantify the achievable resolution for a realistic astrophysical situation, we performed simulations with identical b=20 km/s (Doppler width, Penton et al. 2000) absorption features of variable separation at 1200Å. We define the resolution R80 as the wavelength divided by the separation between identical Gaussian absorption features at which the central maxima between the two minima has a flux deficit of 80% of the feature minima. Sample simulations for two separations are shown in Figure 2.

We find the resolution (R80) of STIS+G140M at 1200Å to be R80≈5800 ($\Delta\lambda$=0.21Å) for the $52 \times 0.1''$ aperture and R80≈4500 ($\Delta\lambda$=0.27Å) for $52 \times 0.2''$. The 'wingless' $52 \times 0.1''$ replacement can resolve features at R80≈6500 ($\Delta\lambda$=0.18Å), and the delta-function replacement can resolve features at R80≈7700 ($\Delta\lambda$=0.16Å) (see Figure 3). This implies that we are able to increase STIS/G140M/$52\times0.2''$ archive spectra from R80≈4500 to 7700; a 71% improvement for our delta-function replacement, and a 46% improvement using the

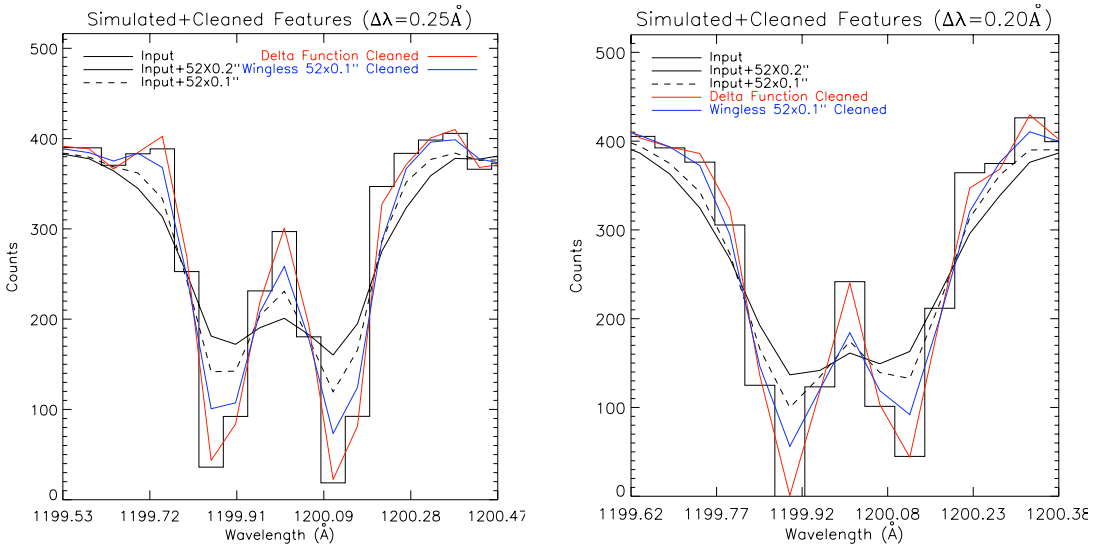

Figure 2: Left: A simulated S/N=20 per pixel STIS/G140M spectrum (rectangular contours) with b=20 km/s absorption features separated by 0.25Å (ΔV$\sim$60 km/s) is convolved with the LSF of the $52 \times 0.2''$ aperture (minima $\sim$170), then cleaned by two different methods (delta function: minima $\sim$40, wingless 52x0.1: minima $\sim$100). Shown for comparison is the convolution of the input spectrum with the $52 \times 0.1''$ LSF (dashed line). Both cleanings of the input $52 \times 0.2''$ spectra achieve improved resolution and are much better estimates of the known input spectrum. By our R80 definition, the input+$52 \times 0.2''$ simulation does not resolve the features while the input+52x0.1$''$ and the cleaned spectra do. Right: For simulated/deconvolved spectral features separated by 0.2Å (ΔV$\sim$50 km/s), neither of the original LSFs resolve the features, yet both spectral cleanings resolve the features.

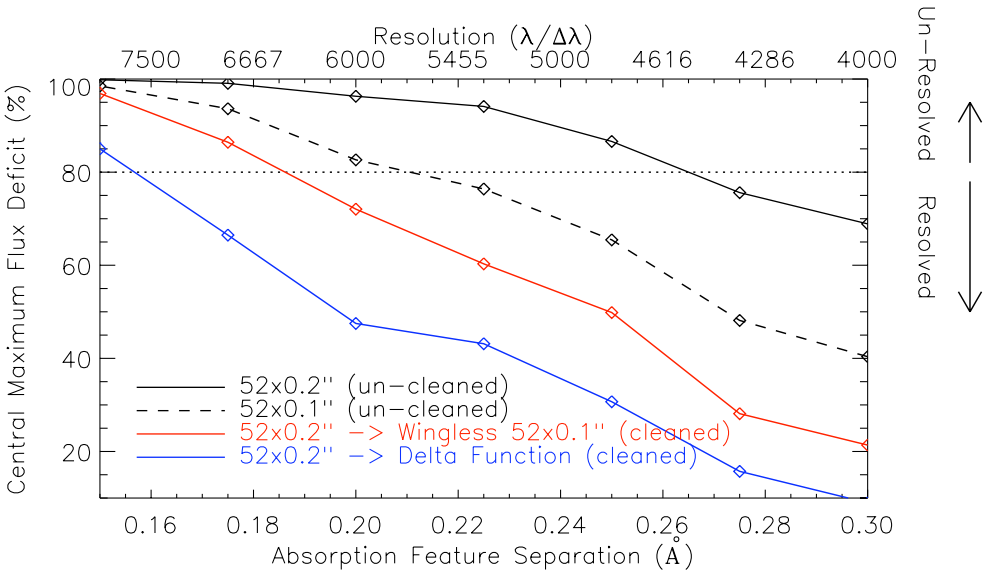

Figure 3: Simulation statistics (see Figure 2) allow us to calculate the relative resolutions of the G140M $52 \times 0.1''$ and $52 \times 0.2''$ apertures (upper solid and dashed lines) compared to spectral cleanings performed on simulated STIS/G140M/$52 \times 0.2''$ data (lower solid lines). Our delta-function and 'wingless' $52 \times 0.1''$ "cleanings" of input+$52 \times 0.2''$ simulations achieve $\sim$70% and $\sim$50% higher spectral resolution than the input+$52 \times 0.2''$ and input+$52 \times 0.1''$ simulations, respectively.

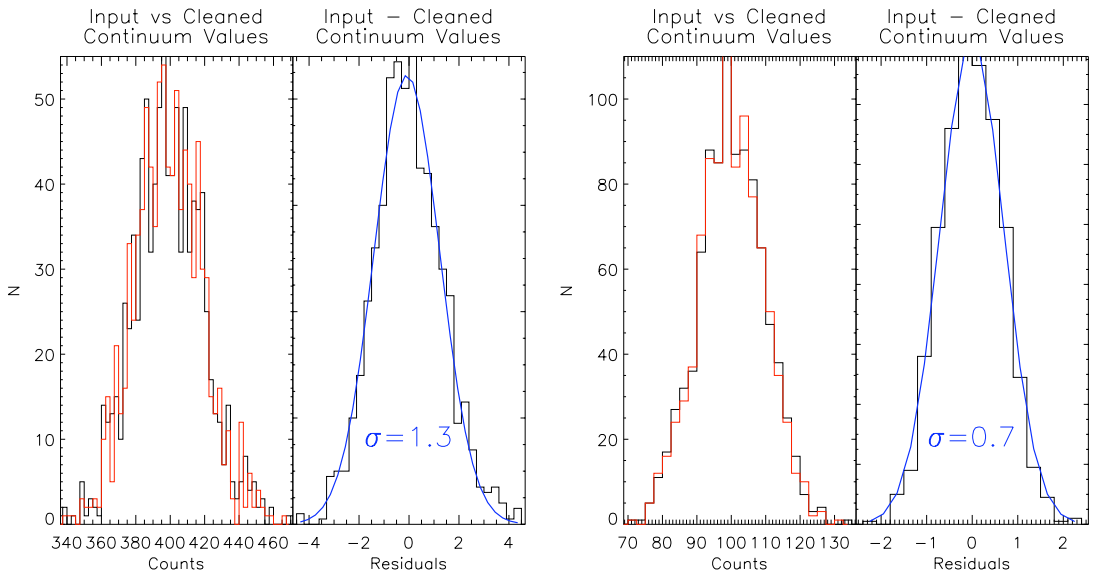

Figure 4: The left panel of the left figure shows the incident (black or dark) and delta-function-cleaned (red or light) continuum values for a simulated G140M 200 Å continuum region ($52 \times 0.2''$ aperture) with S/N per pixel of 20. The right panel of the left figure shows the pixel-to-pixel variations (residuals, black bar graph) between the input and cleaned continuum values. A Gaussian fit of the residuals (blue curve) demonstrates that the residuals appear mainly Gaussian with a σ of ± 1.3 counts or 0.3% (1.3/400). The right figure is identical, except that the S/N=10 per pixel, and the σ is $\pm$ 0.7 counts or 0.7% (0.7/100). In both examples, non-Gaussian excess residuals occur in $< 4\%$ of the cleaned continuum, and the excess is a small fraction of the intrinsic photon noise.

'wingless' $52 \times 0.1''$ replacement. We find similar results for low-resolution STIS spectra, but less dramatic results for echelle data, where the wings of the LSF are less pronounced.

4. Marginal (1%) Increase in Continuum Noise

To examine the effects of our spectral cleaning on continuum noise, we "cleaned" a simulated S/N of 20 (per pixel) STIS/G140M spectrum ($52 \times 0.2''$ aperture) over a continuum of 200Å (4000 pixels, ∼2600 resolution elements). The input spectrum, containing photon noise only, was convolved with the $52 \times 0.2''$ aperture LSF, then "cleaned" using a delta-function replacement. Figure 4 displays a histogram of the before and after continuum counts (left) and the differences (residuals, right) for a S/N of 20 simulated spectrum. We are able to retrieve the input spectrum with a 1-σ error of $\pm$ 1.3 counts, or $1.3/400 = 0.3\%$. For a S/N of 10 spectrum (right panel), we are able to retrieve the continuum to within 0.7%. In both examples, the non-Gaussian component of the residuals is $< 4\%$ and is a small fraction (20%) of the intrinsic photon noise. Further optimization of our algorithm will only reduce this non-Gaussian noise component. The main point is that even using the most extreme output LSF, a delta-function, the added non-Gaussian noise is quite small. This is, of course, a major concern since these non-Gaussian residuals could be misinterpreted as larger σ deviations and misidentified as real absorption or emission features.

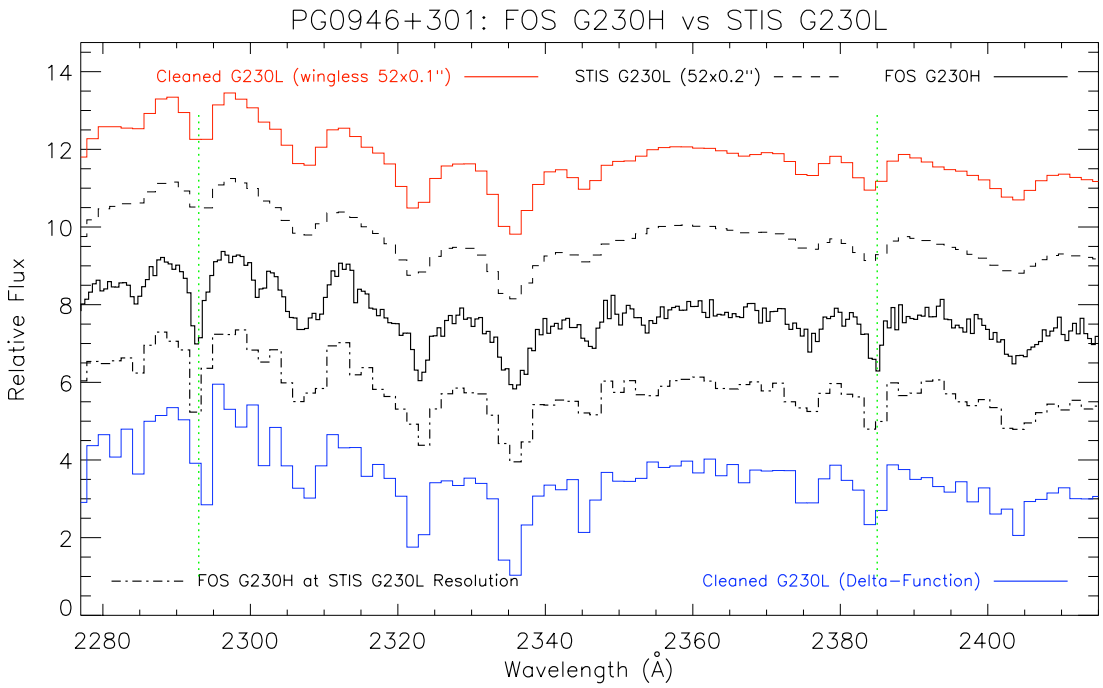

Figure 5: An HST/FOS spectrum of PG 0946+301 (G270H, central solid line) and a pipeline reduction of a similar resolution (G230L STIS, dashed line). Narrow absorption features Lyα at 2293Å and 2383Å are poorly reproduced by the normal STIS pipeline data reduction. Our spectral cleaning algorithm recovers the full extent of the absorption lines. The upper offset (red, top) spectrum is our 'wingless' $52 \times 0.1''$ LSF replacement, while the lower offset (blue, bottom) spectrum is our delta-function replacement. The dot-dashed line shows the FOS/G270H spectrum degraded to the STIS/G230L resolution.

5. HST Example

To demonstrate the operation of this algorithm we compare two spectral cleanings of HST/STIS/G230L/$52 \times 0.2''$ observations of the QSO PG 0946+301 ($z = 1.216$) to the original spectrum and an HST/FOS/G270H spectrum (Figure 5). The two cleaned spectra (the 'wingless' $52 \times 0.1''$ and delta-function replacements) show deeper absorption features, more like the HST/FOS/G270H spectrum. In particular, narrow absorption features, such as Lyα at 2293Å and 2383Å, are poorly reproduced by the normal STIS pipeline data reduction, but are recovered by our spectral cleaning, which looks very similar to the FOS/G270H spectrum degraded to the STIS/G230L resolution.

6. Combining Spectra

At some level, all HST archival spectra should benefit from our spectral cleaning. However, of the 20,400 first-order STIS spectra in the archive, the ~15,600 (76%) that were taken with larger ($> 0.1''$) apertures will benefit the most from our spectral cleaning due to the significant LSF wings. Once "cleaned, these larger aperture datasets can be combined with other "cleaned" observations of the same target taken through other apertures without spectral degradation to produce higher signal-to-noise. In other words, this method allows archive users to place observations taken through different apertures onto a common frame. In addition, the negative impact of non-Gaussian wings of the LSF can be removed, resulting

in increased resolution with minimal contribution to continuum noise. Our algorithm works for all HST modes for which the LSF is well known.

7. Conclusions

By removing the non-Gaussian wings of the LSF, our algorithm increases the resolution of HST archive spectra by as much as 70% for the STIS medium resolution gratings, while introducing $< 1\%$ noise to the continuum. Our algorithm also allows spectra taken through different apertures to be combined in a way that actually increases the S/N and spectral resolution.

7.1. Caveats

- The algorithm only works if the LSF is well-known.
- This algorithm is ONLY designed to correct non-Gaussian LSF wings. If the intrinsic LSF does contain significant wings (i.e., HST/STIS echelle data), then this algorithm will not significantly improve the data, except in the case of combining data taken with different apertures.
- The algorithm is still under development and has not yet been applied to more than a couple of HST archive datasets. So, for now, one must consider this a promising, but unproven, calibration tool.
- There are several subtle nuances of the algorithm (which do not affect the statistics) that are not revealed here due to space limitations.
- Objects not centered in the aperture may have distorted LSFs, which cannot be corrected. This will result in decreased resolution improvement and increased continuum noise.

References

Hőgbom, J. A. 1974, A&AS, 15, 417

Kim Quijano, J., et al. 2004, *STIS Instrument Handbook* (Baltimore: STScI)

Penton, S. V., Stocke, J. T., & Shull, J. M. 2000, ApJS, 130, 121

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Correcting STIS CCD Point-Source Spectra for CTE Loss[1]

Paul Goudfrooij, Ralph C. Bohlin, and Jesús Maíz-Apellániz[2]

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. We review the on-orbit spectroscopic observations that are being used to characterize the Charge Transfer Efficiency (CTE) of the STIS CCD in spectroscopic mode. We parametrize the CTE-related loss for spectrophotometry of point sources in terms of dependencies on the brightness of the source, the background level, the signal in the PSF *outside the standard extraction box*, and the time of observation. Primary constraints on our correction algorithm are provided by measurements of the CTE loss rates for simulated spectra (images of a tungsten lamp taken through slits oriented along the dispersion axis) combined with estimates of CTE losses for actual spectra of spectrophotometric standard stars in the first order CCD modes. For point-source spectra at the standard reference position at the CCD center, CTE losses as large as 30% are corrected to within $\sim 1\%$ RMS after application of the algorithm presented here, rendering the Poisson noise associated with the source detection itself to be the dominant contributor to the total flux calibration uncertainty.

1. Introduction

Since the installation of the Space Telescope Imaging Spectrograph (STIS) onto HST in February 1997, radiation damage to its CCD (which is primarily due to high-energy protons which are especially abundant when crossing the South Atlantic Anomaly) has caused a degradation of its Charge Transfer Efficiency (CTE, defined as the fraction of charge transferred from one pixel to the next during readout). In characterizing the effect of the radiation damage to CCD performance it is often more useful to use the term Charge Transfer *Inefficiency* (CTI = 1−CTE). The observational effect of CTI is that an object whose induced charge has to traverse many pixels before being read out appears to be fainter than the same object observed near the read-out amplifier.

Earlier on-orbit characterizations of the CTI of the STIS CCD have mostly concentrated on the (time-dependent) effect on imaging photometry of point sources of varying signal and background levels as well as measurement aperture sizes (Gilliland, Goudfrooij, & Kimble 1999; Kimble, Goudfrooij, & Gilliland 2000; Goudfrooij & Kimble 2003). However, the CTI corrections reported in those papers are not applicable to spectroscopic observations since the charge structure of spectral data is significantly different from that of imaging exposures. This is illustrated in Figure 1. Consider STIS observations, where the spectral dispersion direction is along rows and the parallel readout direction[3] is along columns. Then, for a given signal level in a measurement element, imaging data features a significantly different signal in each column, whereas spectroscopic data is much more constant along rows. As

[1]Based on observations with the NASA/ESA *Hubble Space Telescope*, obtained at the Space Telescope Science Institute, which is operated by AURA, Inc., under NASA contract NAS5-26555

[2]Affiliated with the Space Telescope Division, European Space Agency

[3]The CTI of the STIS CCD is only significant in the parallel readout direction (e.g., Kimble et al. 2000). Hence, this paper only addresses *parallel* CTI.

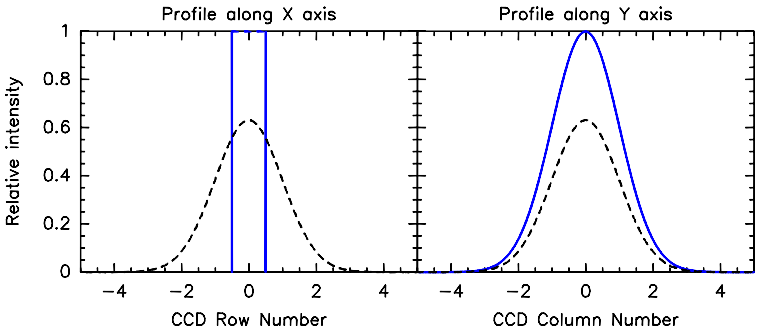

Figure 1: Illustration of the difference in charge structure between imaging and spectroscopic data. The dispersion direction for the latter is assumed to be along rows. For a given signal level per measurement element (per pixel along the dispersion for spectra, per object within a circular aperture for imaging data), spectroscopic data (solid lines) have a higher signal level per column than imaging data (dashed lines). Since the STIS CCD virtually only suffers from CTI in the parallel read-out direction (along the Y axis), this results in a lower CTI for spectroscopic vs. imaging data of a given total measured signal level.

the CTI is highly dependent on signal level and the ratio between background and signal level (see below), this leads to significantly different CTI values for a given total signal.

The purpose of the current paper is to characterize the CTI of the STIS CCD for point-source spectrophotometry in terms of its dependencies on signal level within the spectrum extraction box, background counts outside the extraction box, and elapsed on-orbit time. The results described in this paper supersede those reported in an earlier STIS Instrument Science Report (Bohlin & Goudfrooij 2003, hereafter BG03).

The STIS CCD is a 1024×1024 pixel, backside-illuminated device with $21\,\mu\mathrm{m} \times 21\,\mu\mathrm{m}$ pixels. It was fabricated by Scientific Imaging Technology (SITe) with a coating process that allows it to cover the $200-1000$ nm wavelength range for STIS in a wide variety of imaging and spectroscopic modes. Two serial registers are available. Read-out amplifiers are located at all four corners, each with an independent signal processing chain. By default, science exposures employ full-frame readout through amplifier 'D' (located at the top right of the CCD), which features the lowest read-out noise. Further technical details regarding the STIS CCD are provided in Kimble et al. (1994, 2000).

This paper is organized as follows. Section 2 describes the method used to derive the time constant of the CTI of the STIS CCD in spectroscopic mode. We derive functional dependencies of the spectroscopic CTI on source counts, background counts, and spatial extent of the point-spread function (PSF) in Section 3. We end with concluding remarks.

2. The Time Constant of the CTI of the STIS CCD

We derive the time constant of the CTE degradation of the STIS CCD using a method designated the "Internal Sparse Field" test, which provides measurements that are directly applicable to spectroscopic observations with the STIS CCD. This test was developed by the STIS Instrument Definition Team during ground calibration, and it has been used throughout the on-orbit lifetime of STIS to allow accurate monitoring. It quantifies two key aspects of CTE effects on spectroscopic measurements: *(i)* The amount of charge lost *outside* a standard extraction aperture, and *(ii)* the amount of centroid shift experienced

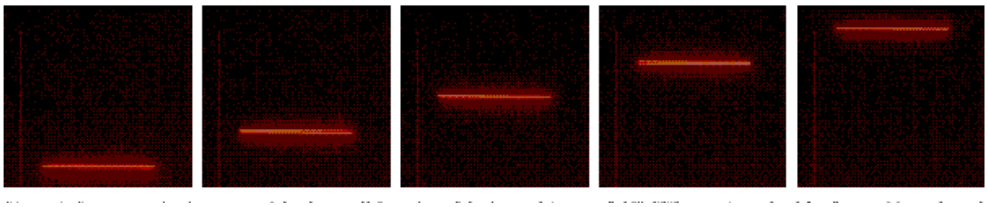

Figure 2: Representative images used for the parallel version of the "Internal Sparse Field" CTE test. A sequence of exposures is taken at each of five positions along the CCD columns, alternating between read-out amplifiers 'D' (located at top right) and 'B' (at bottom right). Systematic variation of the relative intensities observed by the two amplifiers as a function of position reveals the CTE effect (see Fig. 3).

by the charge that remains *within* that extraction aperture. The "internal"[4] version of the sparse field test was implemented as follows. Using an onboard tungsten lamp, the image of a narrow slit is projected at five positions along the CCD columns. At each position, a sequence of exposures is taken, alternating between the 'B' and 'D' amplifiers for readout. This is illustrated in Figure 2. For each exposure, the average flux per column within a 7-row extraction aperture (i.e, the default extraction size for long-slit STIS spectra of point sources, cf. Leitherer & Bohlin 1997) as well as the centroid of the image profile within those 7 rows are calculated. The alternating exposure sequence allows one to separate CTE effects from flux variations produced by warmup of the tungsten lamp. As the slit image extends across hundreds of columns, high statistical precision on CTE values can be obtained even at low signal levels per column. Although these data are not necessarily taken in dispersed mode[5], the illumination is representative for typical spectroscopic observations (as the dispersion direction of STIS CCD spectral modes is along rows). The slit image has a narrow profile (2-pixel FWHM), similar to a point source spectrum.

A key virtue of this method is that neither a correction for flat-field response non-uniformity is required, nor an a-priori knowledge of the source flux (as long as the input source is stable during the alternating exposures). It should be noted that what is being measured is actually a sum of the charge transfer inefficiencies for the two different clocking directions. However, given the identical clocking voltages and waveforms and with the expected symmetry of the radiation damage effects, we believe the assumption that the CTI is equal in the two different directions is a reasonable one.

We emphasize that in calculating CTI from this test, charge is only considered "lost" if it is no longer within the standard 7-row extraction aperture. I.e., we are only measuring the component of CTI produced by relatively long-time-constant charge trapping. Hence, the CTI values derived from this test will not agree with those measured by (e.g.) X-ray stimulation techniques using Fe^{55} or Cd^{109}, for which charge deferred to even the very first trailing pixel formally contributes to the CTI. However, the measurement described here *is* directly relevant to the estimation of CTE effects on STIS spectrophotometry.

2.1. Results

A χ^2-minimization algorithm was used to compute CTI for each observing epoch and signal level. After correcting for (small) gain differences in the two readout amplifier chains, the

[4] "Internal" in this context means that the necessary observations were performed during Earth occultations, hence not requiring any valuable "external" *HST* observing time

[5] About half of the exposures employ STIS in imaging mode using a mirror in the mode select mechanism; the other half use the G430M grating at central wavelength 5471 Å, which produces a very flat spectrum when used with the Tungsten lamp.

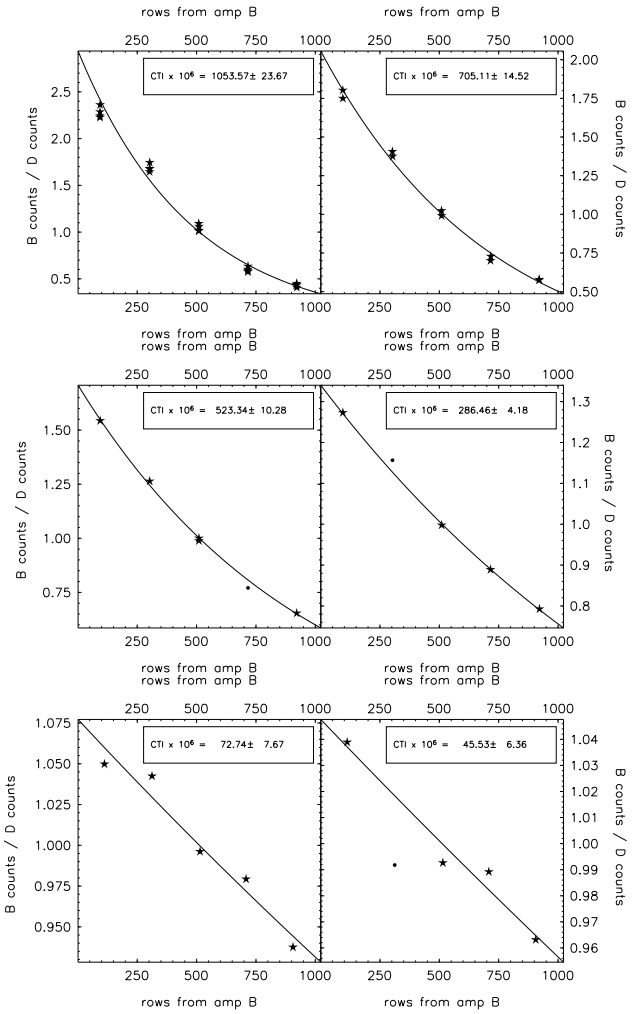

Figure 3: CTI calculation for the internal sparse field measurements of September 2003. Each panel shows the data and the fit for a given signal level. For the panels from the upper left through the upper right, middle left, etc., the signal levels are 60, 130, 195, 500, 3450, and 9850 e^-, respectively. Star symbols indicate measurements used in the fit and circles indicate rejected points. Fitted CTI values are indicated in the boxed legend.

observed ratio of the fluxes measured by the two amps was fit to a simple CTE model of constant fractional charge loss per pixel transfer, allowing for $\kappa - \sigma$ clipping of outliers (the latter arise occasionally from lamp intensity fluctuations of short (0.2 sec) exposures). Flux ratio results for the parallel internal sparse field test taken in `CCDGAIN=1` after 5.5 years in orbit are presented in Figure 3. It can be seen that this simple CTE model fits the data as a function of source position along the parallel clocking direction of the CCD quite well.

To derive the time dependence of the CTI, all CTI measurements were first normalized to a constant background value, i.e. zero background. In order to arrive at that value, two corrections were required: First, the effect of the spurious charge in STIS CCD bias frames (Goudfrooij & Walsh 1997) was accounted for by considering the total background to be the measured one plus the spurious charge. Second, the background dependency of the CTI (to be described in Section 3 below) was taken into account.

CTI values derived as mentioned above for the parallel internal sparse field test taken at different epochs are plotted in Figure 4. In-flight CTE degradation from a pre-flight starting point of low CTI is obvious. Typical CTI behavior is observed as a function of

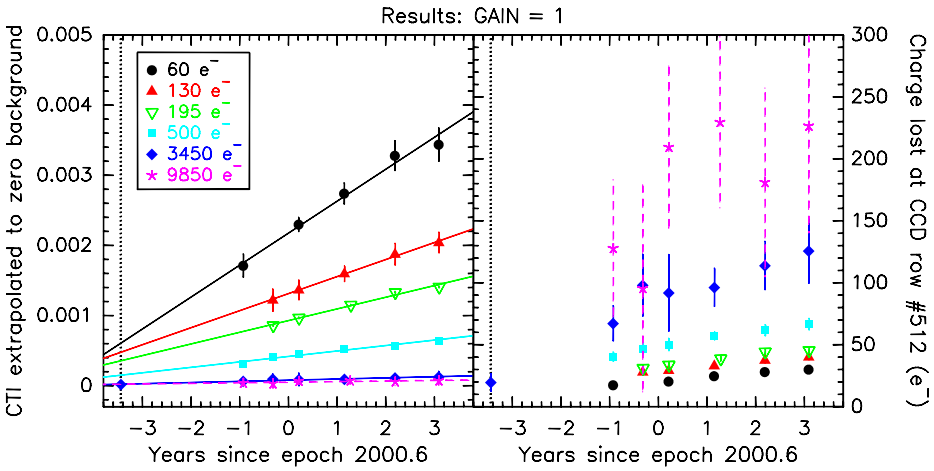

Figure 4: *Left panel:* CTI extrapolated to zero background for `CCDGAIN=1` as a function of time and signal level, derived from the internal sparse field test. Both the data and the corresponding linear fits are plotted. Symbols associated with individual signal levels (corrected for CTI) are indicated in the legend. *Right panel:* Absolute charge lost due to CTI for an object at the central row of the STIS CCD as a function of time and signal level. Symbol types are the same as in the left panel. The epoch of HST Servicing Mission 2 (during which STIS was installed on HST) is depicted as a vertical black dotted line.

signal level: The *fractional* charge loss (which is proportional to CTI) drops with increasing signal level, while the *absolute* level of charge loss increases. The time dependence was derived by fitting the zero-background CTI values to a function of the form:

$$\mathrm{CTI}(t) = \mathrm{CTI}(0)\,[1 + \alpha(t - t_0)], \tag{1}$$

with t in years and $t_0 = 2000.6$, the approximate midpoint in time of in-flight STIS observations.

Results for the time-dependence fit for the `CCDGAIN=1` setting are shown in Fig. 4 and Table 1. The functional fit to the data is quite good, and the derived values for α in Eq. (1) are consistent with one another (within the uncertainties) for all signal levels measured. As to the selection of the time constant α, we considered that the dataset with 3450 electrons per column is the only one for which pre-flight measurements were available, i.e., it covers a time interval considerably longer than for the other signal levels. Hence time constant $\alpha = 0.218 \pm 0.038$ was selected as representative for all signal levels, as indicated in Table 1.

Table 1: CTE degradation time constant α as a function of signal level for `CCDGAIN=1`. The last row lists our adopted value in boldface font.

signal (e^-)	α (yr^{-1})	σ_α (yr^{-1})
60	0.216	0.009
130	0.192	0.013
195	0.188	0.011
500	0.202	0.006
3450	0.218	0.038
9850	0.170	0.052
α = **0.218 ± 0.038**		

3. Insights on CTI From Monitoring of Flux Standard Star Spectra

3.1. Methodology to derive dependence on signal and background levels

The CTI values derived from the internal sparse field test are "worst-case", since there is essentially no background intensity ("sky" or dark current) to provide filling of charge traps in the silicon lattice of the CCD. Hence, additional observations are needed to constrain the functional dependence of CTI on the background and signal levels. For this purpose we build upon the work of BG03 who utilized the large database of spectrophotometric standard star spectra taken on a regular basis (every few months) using a 2 arcsec wide slit. CTI values for spectra of DA white dwarf flux standards GD 71 and LDS 749B taken using the G230LB and G430L gratings are calculated by dividing their measured fluxes by those measured from G230L spectra taken within a few *HST* orbits from one another. The (time-dependent) flux calibration for the G230L mode (which uses the NUV-MAMA detector, which does *not* suffer from CTE loss) is very well established, and accurate to subpercent level (Stys, Bohlin, & Goudfrooij 2004). The standard star spectra used to characterize the CTI effects are listed in Table 2 along with their signal and background levels.

Table 2: List of flux standard star spectra used to characterize CTI effect as function of signal and background level per exposure. All intensities are in e^-.

Rootname	Grating	Flux Standard	Background level	Range in Signal Levels
o6ig10010	G230LB	G191-B2B	0.4	1000 – 5000
o6ig100d0	G750L	G191-B2B	0.5	150 – 7900
o6il01020	G230LB	LDS 749B	1.9	100 – 1800
o8u2200b0	G430L	AGK+81D266	0.5	3000 – 9200
o8v101030	G750L	WD 1657+343	2.5	30 – 750
o8v2040e0	G230LB	GD 71	0.3	120 – 730
o8v204030	G230LB	GD 71	0.1	20 – 170

In evaluating a suitable functional form to characterize the CTI of the STIS CCD in spectroscopic mode, BG03 followed Goudfrooij & Kimble (2003) who showed that the logarithm of CTI scales roughly linearly with the logarithm of the signal level in imaging mode (a glance at panels (b) and (d) in Figure 7 of the current paper shows the same is true for the spectroscopic modes), and that the slope of log (CTI) vs. log (background) decreases systematically with increasing source signal level[6], suggesting a functional form similar to CTI $\propto \beta\, G^{-\gamma}\ \exp(-\delta[B/G]^{\epsilon})$ where B is the background level and G is the gross signal level. After making initial estimates of parameters β through ϵ from bootstrap tests and using the CTI time constant as it was determined in the spring of 2003 ($\alpha = 0.243 \pm 0.042$), BG03 determined a best-fit functional form

$$CTI(B,G,t) = 0.0355\ G^{-0.750}\ (0.243(t-2000.6)+1)\ \exp(-2.97\,(B'/G)^{0.21}) \qquad (2)$$

where B' is the sum of the sky B, the dark current, and the spurious charge (which are all included in G as well). The values for B and G are readily obtained from the output of the X1D routine within CALSTIS to extract 1-D spectra (McGrath, Busko, & Hodge 1999). The efficacy of the correction per Eq. (2) in removing the CTI effect is illustrated in Figure 5. CTI-induced flux errors as high as $\sim$15-20% at low signal levels (~ 100 – 150 e^- per 7-pixel extraction) are reduced to $\lesssim 1.5\%$ by applying Eq. (2), i.e., an improvement of a factor 10.

[6] A likely physical explanation of this effect is that the background charge fills relatively more traps for small charge packets being clocked through than it does for large charge packets.

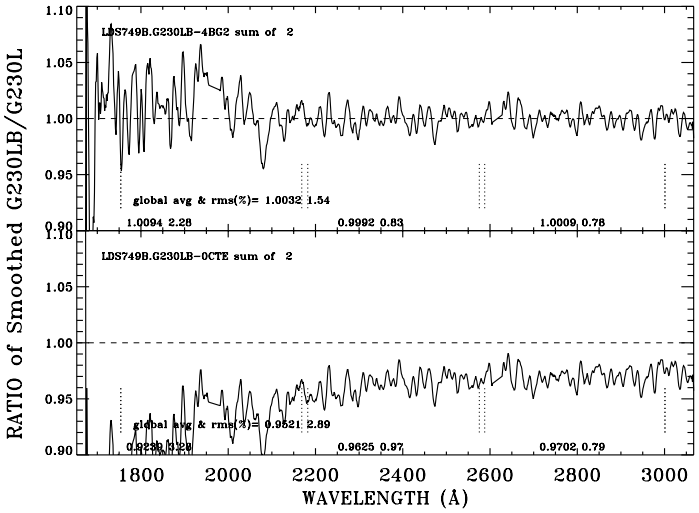

Figure 5: Ratio of the CCD/G230LB flux to the MAMA/G230L flux for LDS 749B. The CCD/G230LB signal level ranges between ~ 90 and 1800 e^-. Both denominator and numerator have been corrected for similar changes of sensitivity with time as per Stys et al. (2004). The bottom panel reveals the G230LB flux error before any CTI correction is applied; the top panel shows the residuals after application of the BG03 CTI algorithm. The global value and rms of the residuals are written on each panel, along with three mean and rms values for the three separate regions delineated by the vertical dashed lines. The bottom panel shows that the error of 8% (0.9229) before CTI correction in the shortest-wavelength region is reduced to $\sim 1\%$ (1.0094) after CTI correction.

3.2. The Impact of the "Red Halo" of the PSF of the STIS CCD

The CTI correction algorithm derived by BG03 was implemented in the CALSTIS pipeline by December 16, 2003, and it still is active at the time of writing. Recently, further testing has shown that application of the BG03 algorithm yields systematic residuals at the red end of the wavelength range covered by the G750L grating. The PSF of the STIS CCD features broad wings at wavelengths $\gtrsim$ 7500 Å (Leitherer & Bohlin 1997), the width of which increases strongly with increasing wavelength. This "red halo" is believed to be due to scatter within the CCD mounting substrate which becomes more pronounced as the silicon transparency increases at long wavelengths. The effects of the red halo are significant, particularly beyond 9500 Å where the default 7-pixel extraction box captures only $\lesssim 70\%$ of the light in the PSF.

This extended halo is likely to have a significant effect on the CTI experienced by the signal within the default 7-pixel extraction box, as the charge induced by the halo acts effectively as 'background' in filling traps. However, the red halo signal is not currently included in the background term (B' in Eq. 2), since the background spectrum used within CALSTIS/X1D is taken far away from the spectrum location[7]. To resolve this issue, we have devised the following update to the CTI correction algorithm of BG03. We split up the background term in two separate terms, B' (as before) and a new term H which contains the fraction of PSF signal *above* the default 7-pixel extraction box. Fortunately, values for H as a function of wavelength can be derived from existing CALSTIS reference files (namely

[7]300 unbinned CCD pixels away by default, as listed in the BK1OFFST and BK2OFFST columns in the 1-D Extraction Parameters Table Reference File, which is listed in data header keyword XTRACTAB.

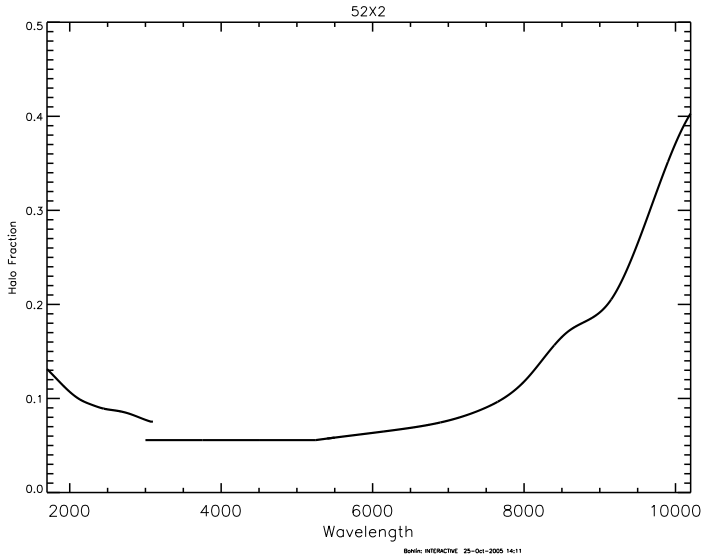

Figure 6: Parameter H in Eq. 4: The fraction of the light in the PSF outside the default 7-pixel extraction box, as a function of wavelength for CCD grating modes. Note the discontinuity near 3000 Å, at the boundary of the wavelength ranges covered by the G230LB and G430L gratings. This is likely due to the presence of a Lyot stop in the G430L and G750L modes which is absent in the G230LB mode (see also contributions of L. Dressel and C. Proffitt in this volume).

from the Photometric Correction Tables[8]). We plot these H values in Figure 6. H is non-negligible at any wavelength, but the spatial extent of the PSF beyond the default extraction box is only a few CCD pixels below $\sim$8000 Å. Hence, low values of H do not necessarily lower the CTI significantly. This effect is accounted for by subtracting a certain minimum threshold value from the measured value of H (i.e., parameter η below).

Considering all of the above, the new functional form of the CTI algorithm is

$$\begin{aligned} \mathrm{CTI} &= (\alpha(t-2000.6)+1)\,\beta\,G^{-\gamma}\ \exp(-\delta[(B'+\epsilon H')/G]^{\zeta}) \\ \text{where } H' &= \max(0.0,(H-\eta))\times Net. \end{aligned} \tag{3}$$

($Net = G-7B$, the net counts in the spectrum.) Initial estimates of the values of parameters β through η and their uncertainties were made using bootstrap tests. A robust fit parameter was then minimized using a non-linear minimization routine from *Numerical Recipes* (Press et al. 1992). Best-fit values of the parameters β through η are listed in Table 3.

The quality of this parametrization of the CTI correction is illustrated in Figure 7, along with a comparison to the BG03 algorithm. Comparing panels (c) and (e) in particular, it is clear that the new solution yields a significantly better correction for the red end of G750L spectra, especially for that of the faint white dwarf WD1657+343 (rootname `o8v101030`). Application of the new CTI correction also renders the STIS G750L fluxes of all faint standards used to determine the apparent NICMOS non-linearity (Bohlin et al. 2005; de Jong et al. 2006, this volume) to be consistent with the relation depicted in Fig. 3 of Bohlin et al. (2005), which lends support to the correctness of the new algorithm. The only spectrum tested for which the new solution still yields a significant residual ($\sim$1%) is the short (35 s) G750L exposure of the DA0 white dwarf G191-B2B (rootname `o6ig100d0`). However, that spectrum shows a 1% error for *any* CTI model at the bright, blue end of the G750L

[8] `*_pct.fits`, listed in data header keyword PCTAB.

Table 3: Best-fit Values of Coefficients in CTI Functional Form (Eq. 4).

Coefficient	Value	Description
α	0.218 ± 0.038	Time dependence of CTI
β	0.056 ± 0.001	CTI normalization
γ	0.82 ± 0.01	Gross count level dependence
δ	3.00 ± 0.05	Normalization for 'background'/gross count ratio
ϵ	1.30 ± 0.10	Multiplicative factor for halo light fraction
ζ	0.18 ± 0.01	Power of 'background'/gross count ratio
η	0.06 ± 0.01	Minimum value of halo light fraction above spectrum

wavelength range, which is unlikely to be due to CTI effects since the CTI algorithm works very well for the other spectra at the same signal and background levels. Furthermore, the new CTI algorithm renders this 1% error to be virtually independent of wavelength for G191-B2B, which again supports the algorithm's correctness.

Overall, the new CTI parametrization formula yields a correction that is accurate within 5% for any data point, while the RMS accuracy for all spectra used in this study stays within 1%. For reference, the dotted lines in panel (e) in Figure 7 show the Poisson noise associated with a resolution element (assumed to be 2 pixels along the dispersion) in spectra of a given signal level. *The new CTI correction formula renders flux calibration to an accuracy better than the uncertainty due to Poisson noise.*

4. Concluding Remarks

We have reviewed the methods used for empirical characterization of the CTI of the STIS CCD and its evolution, using both internal and external exposures which provide measures that are directly applicable to typical spectroscopic observations with the STIS CCD. We derived a functional form for the CTI correction in a semi-empirical fashion, taking into account dependencies on signal level, background level, and the charge trap filling effect of the extended halo in the PSF of the STIS CCD redward of ~ 7500 Å. After applying this CTI correction formula to observed data, systematic residuals stay within 1% (RMS).

The revised CTI correction algorithm presented in this paper will be implemented within the CALSTIS/OTFR pipeline by the next applicable OTFR build. As always, researchers using STIS will be kept in touch on STIS calibration updates like this by email, through the *Space Telescope Analysis Newsletter* which is also available through the "Document Archive" section of the STIS website at `http://www.stsci.edu/hst/stis`.

Acknowledgments. We acknowledge useful discussions with Paul Bristow and invite readers to check out his paper in this volume (Bristow, Kerber, and Rosa 2006), presenting a more physically based CTE correction method for the STIS CCD. We thank Rossy Diaz-Miller for her help in doing tests of the Bristow method and comparisons of its results with those of the empirical method described here.

References

Bohlin, R. C., & Goudfrooij, P., 2003, *Instrument Science Report* STIS 2003-03 (Baltimore: STScI) (BG03), available through http://www.stsci.edu/hst/stis

Bohlin, R. C, Lindler, D. J., & Riess, A., 2005, *Instrument Science Report* NICMOS 2005-02 (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Bristow, P., Kerber, F., Rosa, M.R., 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume

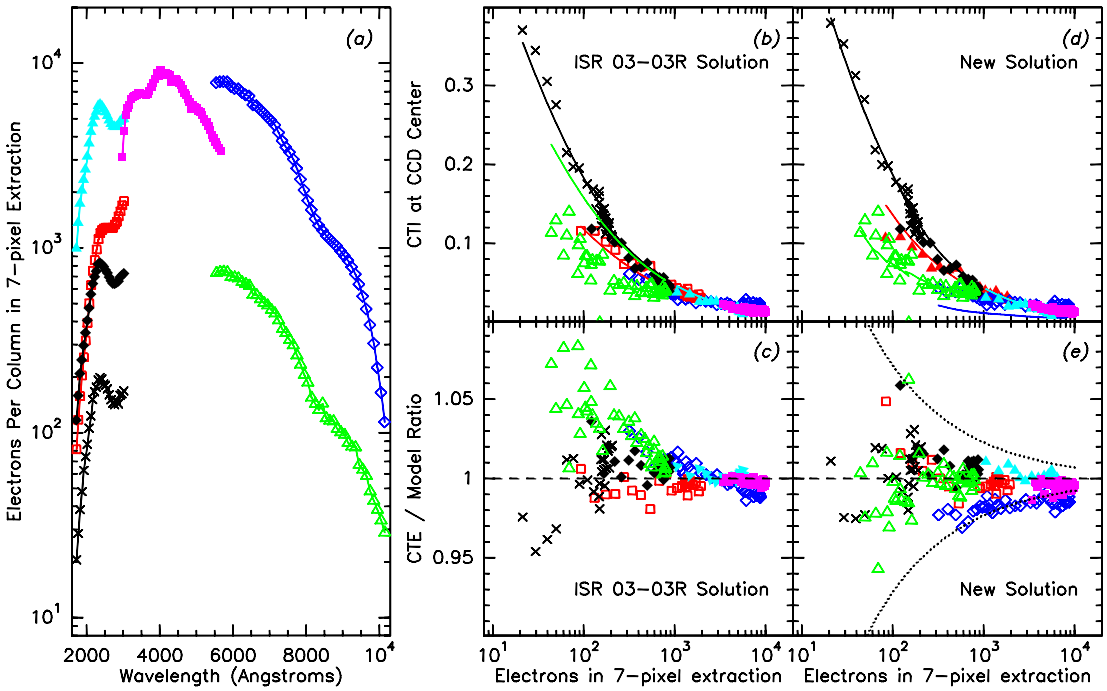

Figure 7: *Panel (a):* Smoothed flux standard star spectra used to determine the functional form of the CTI of the STIS CCD in spectroscopic mode. *Panel (b):* CTI at the central row of the CCD vs. gross signal level within the default 7-pixel extraction box. Symbols represent measured CTI values for the spectra shown in panel (a) using the CTI time constant as determined by BG03 at the time. The drawn lines represent the predictions of the model by BG03 for those data. Symbol types are the same as in panel (a). *Panel (c):* The ratio of measured CTE values and the model predictions by BG03 vs. gross signal level. *Panel (d):* Same as panel (b), but using the CTI time constant and functional form determined in this paper (i.e., Eq. 4). *Panel (e):* Same as panel (c), but using the CTI time constant and functional form determined in this paper. For reference, the dotted lines delineate the uncertainty due to Poisson noise associated with a resolution element of a spectrum with the given signal level.

Gilliland, R. L., Goudfrooij, P., & Kimble, R. A., 1999, PASP, 111, 1009

Goudfrooij, P., & Walsh, J. R., 1997, *Instrument Science Report* STIS 1997-09 (Baltimore: STScI)

Goudfrooij, P., & Kimble, R. A., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 105

Kimble, R. A., Brown, L., Fowler, W. B., Woodgate, B. E., Yagelowich, J. J., et al., 1994, Proc. SPIE, 2282, p. 169

Kimble, R. A., Goudfrooij, P., & Gilliland, R. L., 2000, Proc. SPIE, 4013, p. 532

Leitherer, C., & Bohlin, R. C., 1997, *Instrument Science Report* STIS 97-13 (Baltimore: STScI)

McGrath, M. A., Busko, I., & Hodge, P., 1999, *Instrument Science Report* STIS 1999-03 (Baltimore: STScI)

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T., 1992, Numerical Recipes in Fortran (Cambridge: Cambridge University Press)

Stys, D. J., Bohlin, R. C., & Goudfrooij, P., 2004, *Instrument Science Report* STIS 2004-04 (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

STIS Calibration Enhancement: Wavelength Calibration and CTI

P. Bristow, F. Kerber and M. Rosa

Space Telescope European Co-ordinating Facility, ESO, Karl-Schwarzschildstr 2, 85748, München, Germany

Abstract. The Instrument Physical Modeling Group in the Space Telescope European Co-ordinating Facility has undertaken to examine aspects of STIS calibration and develop physical model based solutions where appropriate. Here we discuss two such aspects and the status of the model based solutions. The first concerns the derivation of dispersion solutions for the echelle spectroscopy modes. The second corrects two dimensional CCD data (imaging and spectroscopic) for the effect of imperfect charge transfer in the read out process.

1. Introduction

The STECF Calibration Enhancement effort for the Space Telescope Imaging Spectrograph (STIS-CE) aims to improve data calibration via the application of physical modeling techniques. A high quality data reduction pipeline for STIS, known as CALSTIS, has been developed by, and is maintained by, the Space Telescope Science Institute (STScI) and the instrument definition team at Goddard Space Flight Center (GSFC), see for example Dressel and Davies (2004) and Proffitt et al (2002) and references therein. This provides users with a rigorous, thorough and reliable calibration based largely on empirical techniques.

On the other hand, STIS-CE concentrates on some aspects of the pipeline where physical modeling may offer some improvement over empirical fits. The performance of existing components of the pipeline provide the obvious benchmark for the success of alternative techniques which we have developed. Hence, in what follows we will often make comparisons between the results of STIS-CE and those of CALSTIS. Whilst we are able to claim some significant improvements, it is testimony to the quality of CALSTIS that it has been quite a challenge to achieve these improvements even for the tasks which we considered most suited to a modeling approach.

Here we focus on two quite different calibration tasks: deriving the dispersion relation for echelle modes and correcting for charge transfer effects in the CCD readout. In both cases we begin by discussing the existing empirical approach. We do this for several reasons. Firstly this helps to introduce the calibration task we are trying to perform and provides a context for our description of the modeling approach. Secondly a description of the empirical approach is helpful for the comparisons that we will make later. Finally, since we began STIS-CE after STIS was already in orbit and the calibration plan was already well established, we were constrained to develop an approach which would take advantage of existing calibration data.

2. Wavelength Calibration

2.1. Empirical Approach

The empirical approach employed in CALSTIS is described in some detail in McGrath et al (1999). The basic principle can be briefly summarized as follows:

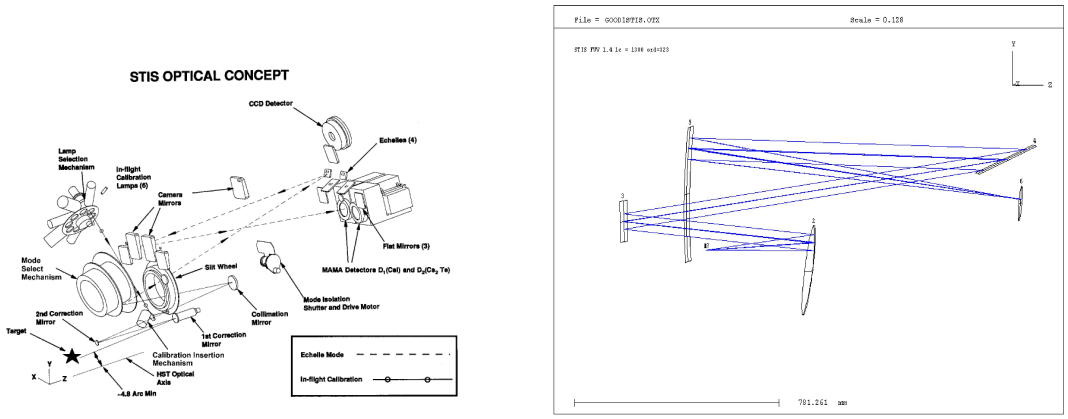

Figure 1: The optical path of the STIS echelle (left) represented as a simple ray trace (right).

1. For each mode:
   - Dedicated deep calibration exposures of the onboard calibration lamps are obtained periodically.
   - Known features in the lamp are identified and a polynomial fit is made which maps pixel positions in a given order to wavelength.
   - The result is stored as a master table for each mode.
2. For science exposures:
   - A short wavelength calibration, hereafter *wavecal*, exposure is made immediately before or after the science exposures.
   - The pipeline reduction finds the appropriate tabulated polynomial solution for the mode.
   - Cross correlation allows the determination of zeropoint offsets between the contemporaneous wavecal and the polynomial solution.

The essential ingredients here are the deep and short wavecal exposures and a reliable list of spectral features for the calibration lamp. Kerber, Bristow, Rosa, Nave et al. (2006) in this volume describe a project which has improved our knowledge of the spectrum emitted by the Pt/Cr-Ne calibration lamps carried by STIS.

2.2. Simple Ray Trace Model

The design of modern astronomical instruments takes advantage of advanced optical modeling software, so we have a detailed knowledge of the optical path. However, this knowledge is not traditionally used in the calibration. In order to replace the polynomial dispersion solution with a physically motivated model, we used our knowledge of the optical components to build a simple ray trace of the STIS optical path, by representing optical surfaces as a series of matrix transformations. In this way were able to compute the eventual detector location of a ray of any wavelength. Of course, this ray trace will describe STIS only if the various tips, tilts, and other properties of the components are all accurately present in the matrices. Furthermore, some of these parameters will be mode dependent and others may change in a less predictable way with time.

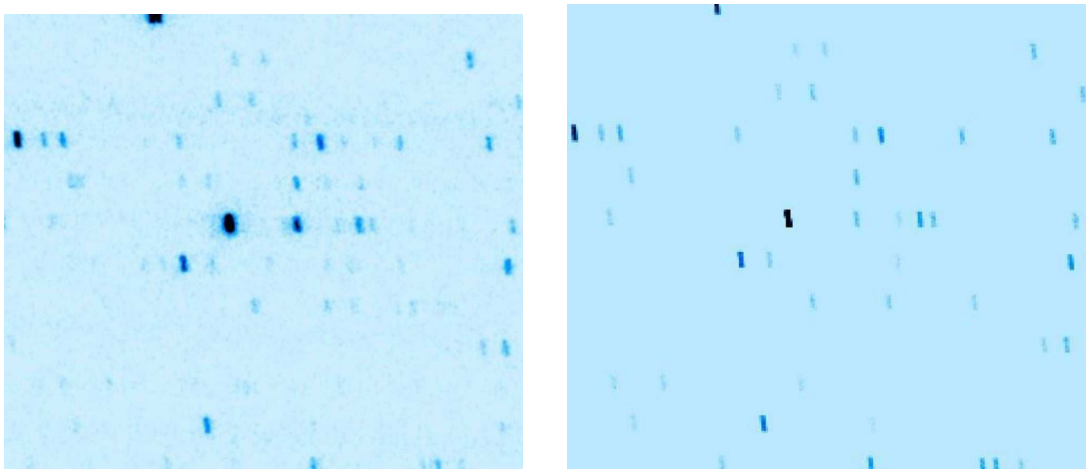

Figure 2: On the left is a section of an exposure of the calibration lamp (E140H), whilst on the right is a simulation produced with the model and a spectral atlas for the lamp.

2.3. Optimization

In order to obtain the precise values of these parameters we go first to the same deep wavecals as used by the empirical approach and require that the model accurately reproduces the distribution of spectral features in the detector plane. However, even our streamlined ray trace of STIS requires no less than 35 parameters. Clearly a trial and error approach is not feasible. Consequently we employ an optimization process known as Simulated Annealing (SA) (Kirkpatrick et al (1983), Press et al (1986)). We provide the optimization routine with a list of wavelengths of known features and their detector locations in the wavecal exposure. The SA algorithm then iteratively calls the optical model while adjusting the parameters and attempting to minimize the difference between the model predicted detector positions for these features and the values measured in the wavecal. The result of this process is a model parameter configuration file which accurately describes the state of the instrument at the time of the deep wavecal exposure.

2.4. Specialization to Short Wavecals

Just as the empirical calibration requires an additional shift in order for the polynomial derived from deep wavecals to fit the contemporaneous short wavecal, the model based approach also makes use of fine tuning. By examining the differences between the optimized configurations obtained for a range of different wavelength settings we identified two parameters (both angles describing the orientation of the cross-disperser) as being critical. Therefore, when fine tuning to the configuration for a given observation a more physical approach than an x,y shift in the detector plane is to adjust these two parameters. This is achieved by applying SA once again, this time requiring that features in the short contemporaneous wavecal are matched and allowing only the two critical parameters to vary.

2.5. Visualization Tools

Once we have determined the configuration of the model components in this way we can use the model to produce simulated data and other "visualization products" with very little further effort. Figure 2 is an example using the calibration lamp line list. Given a model source spectrum and sky emission and absorption spectra, we could also simulate science data and use the model as an advanced exposure time calculator.

Another example is the creation of a map of the wavelengths across the detector generated with the model. In this map the pixel values are the wavelength of light that would

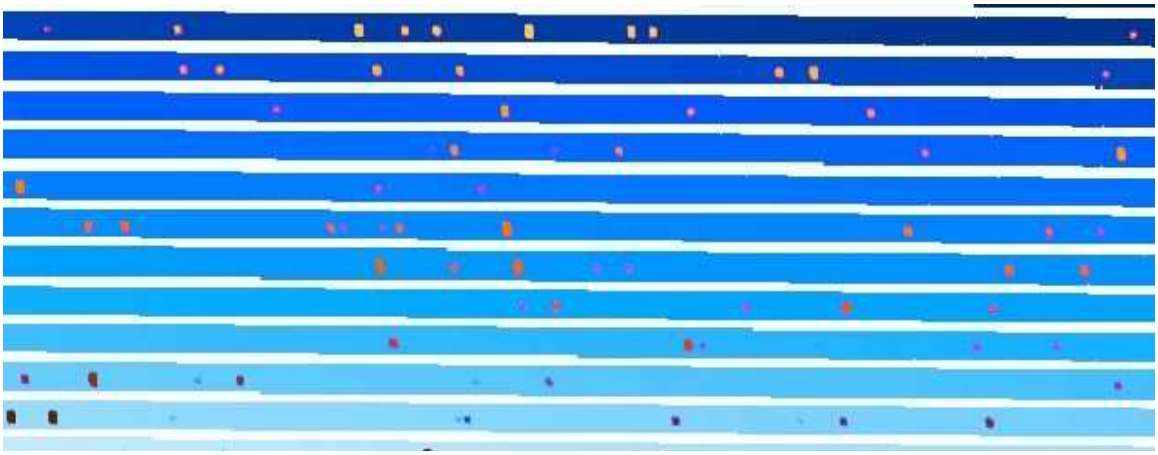

Figure 3: Superposition of a data frame and wavelength map. In practice the simplest way to use such a wavemap would be to blink between the data frame and wavemap in a FITS viewer such as DS9.

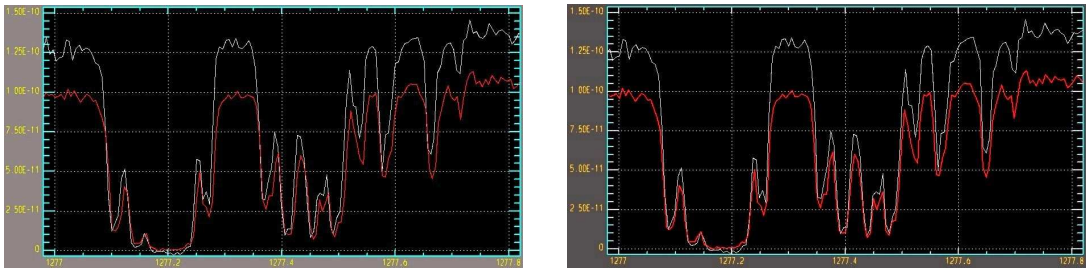

Figure 4: Alignment between features in the overlap regions between two adjacent orders. The upper curve is order 329 and the lower is order 330. a) Empirical calibration. b) Model based calibration.

fall at this detector location. By loading both data and the corresponding wavemap into a FITS viewer such as DS9, pointing at a feature and blinking the images, the user gets a quick estimate of the wavelength of the feature. (See figure 3.)

2.6. Derivation of the Dispersion Relation

The visualisation tools described in the previous section are essentially just spin-offs of this approach. The real advantage is the more accurate dispersion relation that it offers. The ray trace itself is not reversible, so we cannot simply run it backwards to obtain the wavelength corresponding to the center of each pixel. Instead we simply call the model for a list of wavelengths covering the range of the mode in question. These wavelengths are spaced finely enough that the predicted detector positions are separated less than 0.5 pixel. Interpolation is used to establish the wavelength corresponding to the center of each pixel.

2.7. Science Case

As an example of scientific results which may be affected by this correction, we identified a study of Thermal Pressures in the ISM by Jenkins & Tripp (2001). These authors actually noted inconsistencies in the *overlap regions* in their E140H data. The overlap regions are those parts of echelle spectral orders that contain wavelengths also present in previous or subsequent orders. Photons of the same wavelength may reach different locations on the detector, depending which order they are diffracted into. An accurate dispersion solution

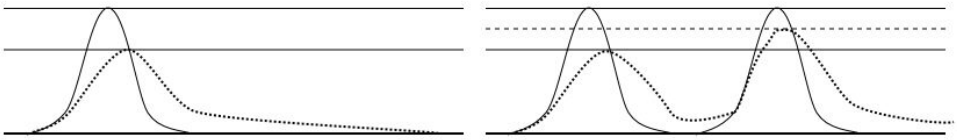

Figure 5: Schematic diagram of the redistribution of charge that occurs during inefficient read out. a) An isolated point source. b) Two sources closely spaced in the readout direction.

should assign the same wavelength to these photons (or rather the pixels they fall into) regardless of which order they are in. What Jenkins and Tripp noticed in their CALSTIS reduced data was a small shift between the wavelengths assigned to spectral features at one end of an order and those assigned to identical features at the beginning of the next (fig. 4a).

We used the model to compute a dispersion solution and found a much better fit (fig. 4b). Further analysis shows that typical offsets in order overlap regions in Jenkins and Tripp's data is of order 0.5 pixels. With the model this is typically reduced to 0.1 pixels. See Kerber, Bristow, and Rosa (2006) in this volume for a more detailed discussion.

3. Charge Transfer Inefficiency (CTI)

By now, inefficient charge transfer is a well known source of photometry errors, degraded point spread function (PSF) and image artifacts (trails) in HST data. The magnitude of the effect increases with the amount of time a CCD detector is on orbit. Bombardment from energetic particles in the hostile radiation environment causes the build up of defects in the silicon lattice which trap charge during the transfer process. Hence WFPC2 data is the worst affected and even ACS data already shows measurable CTI. The removal of charge from large charge packets during readout clearly manifests itself as removal of flux from the central isophotes. Some of this charge is re-emitted in subsequent pixels and, in addition to the effects already mentioned, the centroid is shifted in the opposite direction to the readout.

Figures 5a and b show a grossly exaggerated CTI effect: a) refers to an isolated point source and b) to two sources closely spaced in the readout direction. In both cases the solid curve represents the *true* flux distribution and the dotted line how it would appear after readout with imperfect charge transfer. Also note that the area under the dotted curve is, in reality, always less than that under the solid curve - not all charge is re-emitted. In figure 5b, the source nearer to the register looses charge in the same way as the isolated source in figure 5a. However, in doing so it fills the traps on the CCD in the pixels it is clocked through. As a result, when the charge packets from the second source are transferred through these pixels, many of the traps are already full. This results in less attenuation of charge from the second source. The tail following the second source is similar to that which followed the isolated source. These profiles could equally well represent the cross section (in the readout direction) of a stellar image, or the cross-dispersion profile of a spectrum read out along the slit.

3.1. Empirical Approach

The most dramatic impact upon scientific results is in the form of distorted photometric results from both imaging and spectroscopic data. This can be, to some extent, amelio-

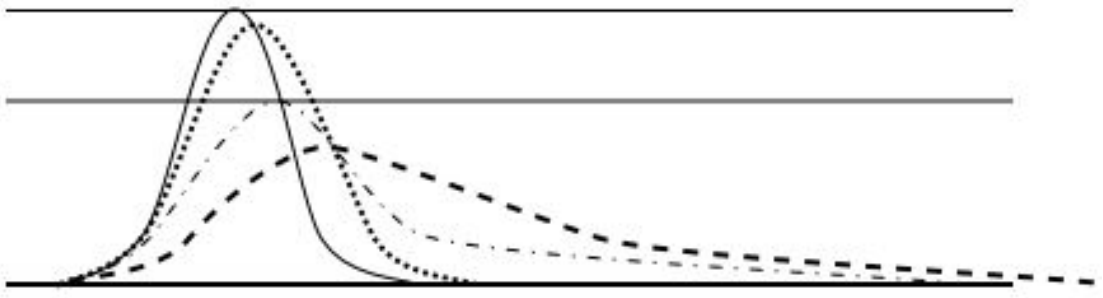

Figure 6: Correction process (see text) illustrated for a point source. Once again, the CTI effects are grossly exaggerated for illustrative purposes

rated by the use of an empirical formula which estimates the flux lost as a function of the position on the chip, the flux of the source, the background level and the level of radiation damage (time on orbit). Goudfrooij & Kimble (2003) and Bohlin & Goudfrooij (2003) have calibrated such functions for STIS imaging and spectroscopic observations respectively.

However, this correction applies only to photometry, not astrometry (centroid shifts) or structure. Moreover, it applies only to the type of objects for which it was calibrated: point sources. It must be calibrated separately for imaging and spectroscopic data, and the calibration requires the acquisition of a considerable amount of calibration data. Finally, this approach cannot account for the situation depicted in figure 5b, which makes it inappropriate in the case of crowded fields or extended objects.

3.2. Readout Model

A physical model based approach, in which CTI effects are understood in terms of the CCD operation and environment, has the potential to consider the flux distribution across the array, thereby returning flux to its true distribution. In addition to the photometry correction provided by the return of flux to central isophotes, the effects upon PSF and centroid and trails are also dealt with. Moreover such a correction applies equally to crowded fields and extended objects as well as point sources. In principle such an approach should be portable to other space based detectors.

We have developed a model (Bristow 2003a) which works in the following way: Starting from two arrays which describe the distribution of charge on a CCD after an exposure and the distribution of traps, we simulate the charge packet shifts that occur during readout keeping track of trapping, emission and dark current. Trap time-scales (as a function of operating temperature) are taken from laboratory data for three different defects known to occur in CCDs subjected to radiation damage. Also critical are the details of the chip architecture, in particular the clocking frequencies, well size and dark current.

3.3. Derivation of Correction

The readout model described is a "forward" model which cannot be run in reverse. So it is in fact only useful for adding CTI effects to data. That is, if we knew the "true" flux distribution (the solid line in figure 6), then the model would help us simulate the CTI degraded flux distribution (the dash-dot line). If we want to remove the effects then we need to proceed as follows:

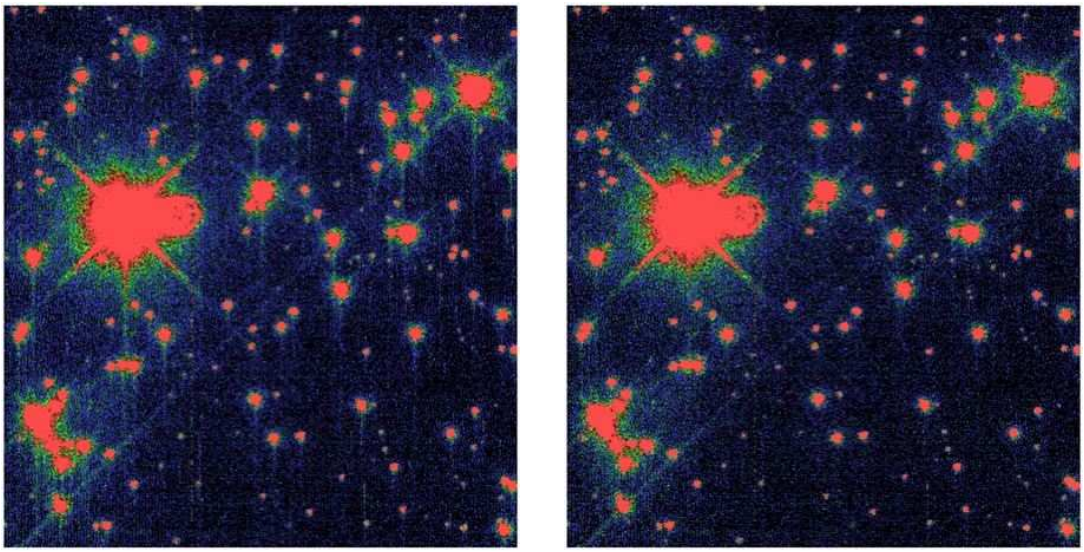

Figure 7: a) Left: Raw STIS CCD imaging data. b) Right: The same data after correction with the readout model.

- Since we have no other "best guess" with which to start the readout simulation, we use raw data (dash-dot line) from an exposure that is to be corrected.
- The simulation then adds simulated CTI to this profile to give the dashed line.
- An improved estimate of the original profile can then be obtained by taking the difference between the dash-dot and dashed profiles and subtracting that from the dash-dot profile to get the dotted corrected profile. Essentially the difference between the dash-dot and dashed profiles represents the charge lost or gained due to CTI.
- Ideally we would run further iterations, taking the dotted corrected profile as input and so on. However, in practice, where the CTI effects are actually much smaller, one iteration is enough to restore 90-95% of the flux to the central isophotes.

This correction must be applied not only to all raw exposures in a dataset, but also to the associated dark and bias frames. This is discussed in detail in Bristow (2004a).

3.4. Validation

The removal of CTI trails is illustrated in figures 7a and b, which show data before and after the application of the model based correction respectively.

In order to see how successfully the simulation derived corrections restore the photometric properties of CTI degraded sources we use the empirical corrections of Goudfrooij & Kimble (2003) as a comparison. These corrections have been calibrated for a range of epochs, background levels, source levels and distances from the readout register for a large number of datasets by using the ability of the STIS CCD to readout in two directions to registers at the top or bottom of the chip. As such they represent a distillation of the available data on CTI, so it is not necessary to repeat the exercises of extracting this information from the data.

Figure 8 represents such a comparison for a large sample of data and shows, on the whole, a good agreement between empirical and simulation derived corrections. There is however some scatter and a systematic overestimate from the simulation derived corrections relative to the empirical corrections as the corrections become larger (lower values in the

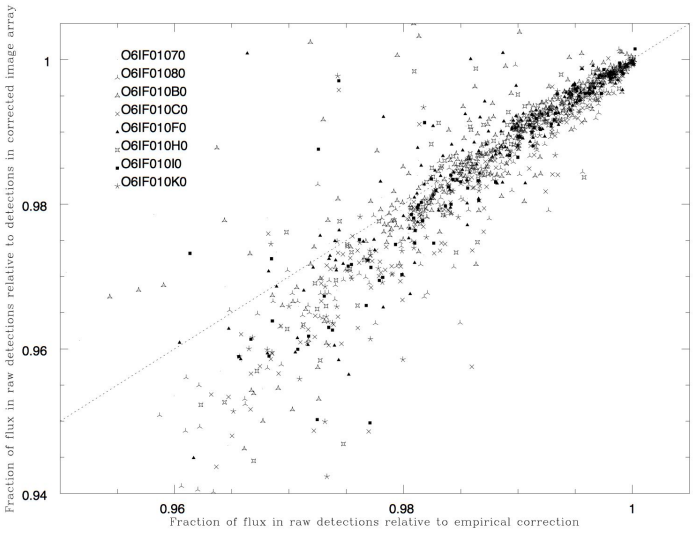

Figure 8: Comparison between model and empirical photometric corrections for the flux within a 7 pixel diameter aperture. Detections come from the datasets according to the key in the figure. The dotted line represents perfect agreement between model and empirical corrections.

plot). The scatter merely reflects the fact that the simulation deals with the actual, inherently noisy, charge distribution. Indeed the scatter represents the failing of the empirical solution for objects which are not bright point sources. If one considers only the bright point sources, as in figure 9, to which the empirical solution applies, then the agreement is much better. Examination of individual sources which fall far from the line of agreement in this plot almost always reveals some aspect (non-gaussian or faint profile or preceding object nearby in the readout direction). These details are discussed in more detail in Bristow 2003b.

3.5. Science Case

The usefulness of this approach is best demonstrated by considering a science case for which astrometry is critical and where the existing empirical corrections have nothing to offer. The example we chose was a measurement of the proper motions of local dwarfs by making observations of constituent stars over a time baseline of several years with STIS and WFPC2 (Piatek et al. 2002 & 2003). Each field was chosen to contain a background QSO as a truly stationary reference point. The analysis in the original AJ papers dealt with issues such as PSFs and geometric distortion in some detail. However the possible shift of the stellar centroids due to CTI was not considered.

To estimate the size of the effect, we corrected a pair of STIS datasets used by Piatek et al (2002, 2003) for CTI and compared the new centroid position for all sources (figure 10). The detector imaged the same position on the sky at the same orientation in both datasets. The crosses, fit by the solid line, are the CTI shifts in the earlier (April 2000) data. The squares, fit by the dotted line, are the CTI shifts in the later (April 2002) data. The CTI shift for the QSO at each epoch is ringed. As expected, a row dependent shift is apparent. If this were the same for each dataset, then the CTI correction would cancel out. However the gradient has steepened in the 2002 data due to the increased radiation damage. See Bristow (2004b) for further discussion.

The exact impact upon the final results, which utilize data from several fields at different inclinations for which the readout direction is different, is very hard to predict in

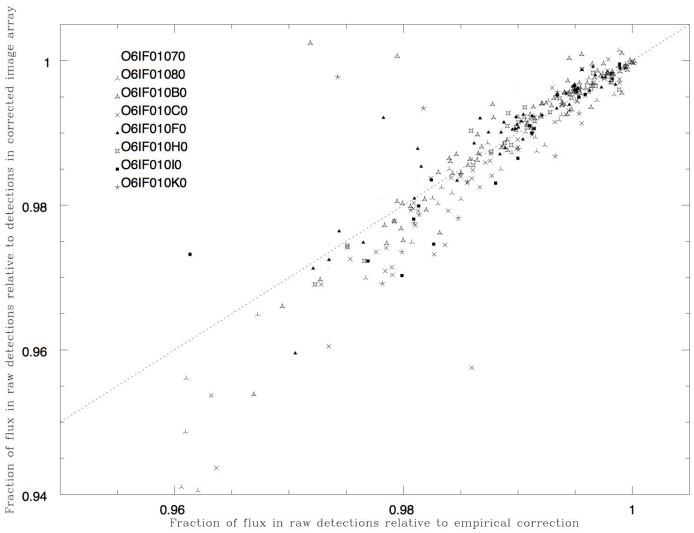

Figure 9: Same as figure 8 except showing only sources for which the empirical correction is supposed to apply.

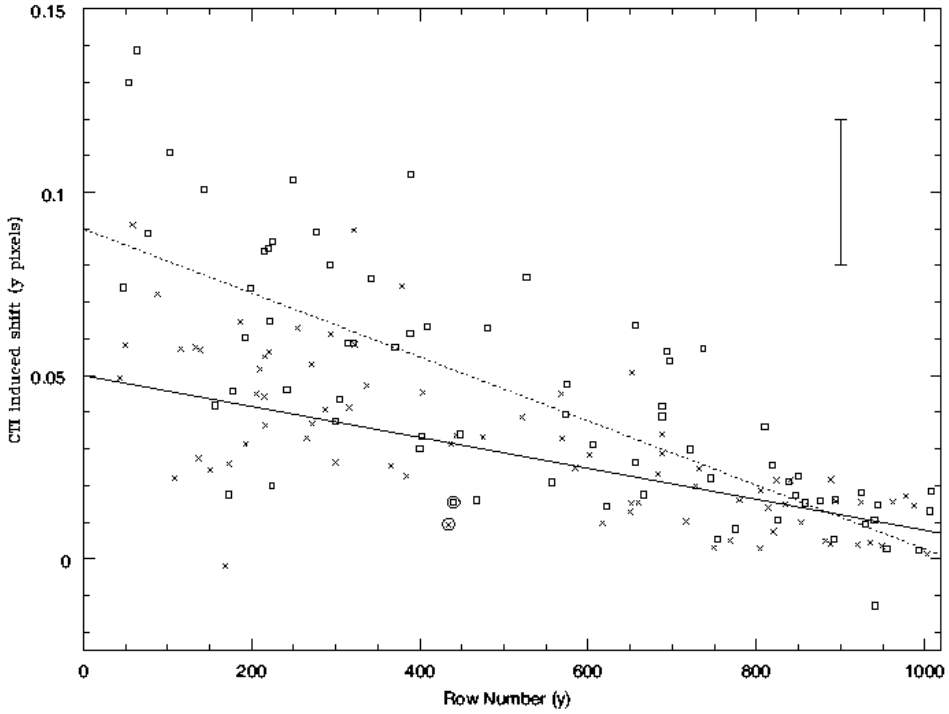

Figure 10: CTI induced centroid y-shift for sources in the earlier epoch O5BL05010 (crosses) and later epoch O6D905040 (squares) datasets. The solid and dotted lines are linear fits to the O5BL05010 and O6D905040 data respectively. The QSO, as it appears in each image, is ringed. For comparison, the error bar represents the scatter in centroid differences for sources from different epochs when put into the same reference frame in Piatek et al (2003).

this way so we provided Piatek et al with corrected data which they submitted to their detailed analysis. The effect upon the derived proper motions was found to be significant and appears to make the results from different fields more homogeneous and reduces the discrepancy between STIS and WFPC2 results (presumably the outstanding disagreement with WFPC2 results is due to CTI in WFPC2 data), Piatek et al (2005). This has led to an ongoing collaboration with data for the dwarf galaxies Ursa Minor (published, Piatek et al 2005), Sculptor (in press) and Carina (processed).

4. Conclusions

We identified aspects of the STIS calibration procedure which lent themselves to a physical modeling approach and where we might expect such an approach to improve upon the already excellent performance of the CALSTIS pipeline. For two such aspects, STIS Echelle wavelength calibration and STIS CCD CTI correction, we developed the necessary physical models and methods of application. These methods have been implemented, tested and applied to science cases yielding significant improvements with respect to the empirical approach.

References

Bohlin, R. and Goudfrooij, P. 2003, *Instrument Science Report* STIS 2003-03 (Baltimore: STScI), available through http://www.stsci.edu/hst/stis

Bristow, P., ST-ECF Newsletter pp9, No. 34, 2003a

Bristow, P., CE-STIS-ISR 2003-001, 2003b

Bristow, P., CE-STIS-ISR 2004-002, 2004a

Bristow, P., CE-STIS-ISR 2004-003, 2004b

Dressel, L. and Davies, J., 2004, *Instrument Science Report* STIS 2004-06 (Baltimore: STScI)

Goudfrooij, P. & Kimble, R. A., 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI)

Jenkins, E.B. & Tripp, T.M., 2001, ApJS 137, p. 297

Kerber, F., Bristow, P., Rosa, M. R., Nave, G., Reader, J., Sansonetti, C. J., & Lercher, G.,2006, this volume, 318

Kerber, F., Bristow, P., & Rosa, M. R., 2006, this volume, 309

Kirkpatrick, S., Gelatt Jr., C.D., & Vecchi, M.P. 1983, Optimization by Simulated Annealing, Science, V. 220, No. 4598, p. 671

McGrath, M. A., Busko, I. & Hodge, P. 1999, *Instrument Science Report* STIS 99-03 (Baltimore: STScI)

Piatek, S., et al. 2002, AJ, 124, 3198

Piatek, S., et al. 2003, AJ, 126, 2346

Piatek, S., et al. 2005, AJ, 130, 95

Press, W.H., Flannery, B.P. Teukolsky, S.A. & Vetterling, W.T., 1986, Numerical Recipes: The Art of Scientific Computing (Cambridge Univ. Press,Cambridge, UK)

C. R. Proffitt, P. Goudfrooij, T. M. Brown, J. E. Davies, R. I. Diaz-Miller, L. Dressel, J. Kim Quijano, J. Maz-Apellniz, B. Mobasher, M. Potter, K. C. Sahu, D. J. Stys, J. Valenti, N. R. Walborn, R. C. Bohlin, P. Barrett, I. Busko, and P. Hodge, 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

STIS Calibration Enhancement (STIS-CE): Dispersion Solutions Based on a Physical Instrument Model

Florian Kerber, Paul Bristow, Michael R. Rosa[1]

Space Telescope European Co-ordinating Facility, Garching, Germany

Abstract.

The Space Telescope European Co-ordinating Facility (ST-ECF) embarked on the STIS calibration enhancement (STIS-CE) effort as part of the extension of the Memorandum of Understanding between NASA and ESA for the Hubble Space Telescope (HST). The work was done by the ST-ECF's Instrument Physical Modeling Group in collaboration with the STScI's Spectrographs group.

Traditionally, calibration of spectrographs is done by empirical methods, e.g. a wavelength dispersion solution is derived by applying a polynomial fit to the emission line spectrum of a calibration exposure. In the model based approach we make full use of the engineering information used to build the instrument. Details of the concept are described in Rosa (1995). Our approach comprises a ray trace model (Ballester & Rosa 1997) describing the geometry of the optical elements and the software needed to compare the performance of the instrument as determined by the model with actual calibration data. We then use an optimizer tool which finds, for any given observation, the actual configuration of the instrument. With this method we can close the loop between model description and actual observations resulting in a high fidelity calibration of the instrument. For STIS we have completed the dispersion solution of the high resolution echelle modes and have demonstrated that all central wavelength settings of E140H can be properly described by changing only two tilt angles of the cross-disperser. An integral part of our approach is the provision of high quality input data: in the case of STIS our work - in collaboration with the Atomic Spectroscopy Group at the US National Institute of Standards and Technology (NIST) - on the output of the calibration lamps has remedied a long-standing shortcoming of the calibration. The standard wavelength calibration of all HST spectrographs has been based on a line list produced by Reader et al. (1990) for a Pt-Ne lamp. STIS uses a Pt/Cr-Ne lamp; the addition of Cr is especially significant in the near UV where up to 90 % of the observed lines are Cr. Without established wavelength standards these could not be used.

Our new dispersion solution based on an instrument model and using the new Pt/Cr-Ne line list (Sansonetti et al. 2004) has been successfully applied to scientific data. In order to investigate the improvement achieved in a quantitative manner we have chosen a science case looking at interstellar absorption lines. For the echelle mode E140H we have compared the wavelengths of the same lines in the overlap regions of adjacent orders. Using the CALSTIS pipeline we find a systematic offset of order 0.5 pixel which is within the expected accuracy. The model based solution delivers a factor of five improvement bringing the difference in wavelength to within order 0.1 pixel. Given the fact that STIS is a very well calibrated instrument this nicely illustrates the potential of the physical modeling approach.

[1]affiliated to the Space Telescope Operations Division of Research and Scientific Support Department of the European Space Agency

1. Introduction

For the past 2.5 years the ST-ECF's Instrument Physical Modeling Group (IPMG) has been working on an enhanced calibration for STIS (STIS-CE). This effort has been an integral part of the extension of the Memorandum of Understanding between NASA and ESA for the Hubble Space Telescope (HST). The work was done in collaboration with the Spectrographs group at the Space Telescope Science Institute (STScI).

STIS currently is arguably the best calibrated UV-visible Echelle Spectrograph in Astronomy. STScI provides science ready products that are the result of a pipeline reduction using a dedicated software CALSTIS. This software is kept up to date using the best available calibration and can also be used by the users in the community in an interactive fashion if certain aspects of the calibration deserve particular attention. The calibration provided uses traditional empirical methods, e.g. a wavelength dispersion solution is derived by applying a 2-D polynomial fit to the emission line spectrum of a calibration exposure.

2. Physical Modeling

For the physical model we use the engineering information used to build the instrument and our understanding of the physics governing the behavior of the optical elements. In this manner we are able to describe how the instrument would perform if exactly built to the specifications but we can also investigate what effects small deviations from the planned instrument will have. Note that this approach is not restricted to the optical performance of the instrument. For STIS we have also developed a physical description of the read out process of the CCDs and charge transfer inefficiency caused by radiation damage. Details of this model are described in Bristow, Kerber & Rosa (2006).
Since the instrument is subject to changes over time caused by, e.g., outgassing of the support structure or environmental influences such as temperature changes along the orbit, the actual configuration of the instrument is time dependent. Of course, during operations – e.g. change of grating or central wavelength – the configuration of the instrument also changes frequently. The model based approach now has the big advantage of being able to describe such variations in terms of actual changes of parameters relating to physical properties of the instrument. For STIS, for example, we have demonstrated that all central wavelength settings of E140H can properly be described by changing only two tilt angles of the cross-disperser.

2.1. Modeling Technique

The physical model is based on the specifications for the optical design of STIS; see Fig. 1a in Bristow, Kerber & Rosa (2006). We used the following ingredients – see Fig. 1 for a schematic view – to describe the instrument and to derive a dispersion solution:

- a ray trace model describing the geometry of the optical elements
- Monte-Carlo illumination of the slit to simulate observations of, e.g., the calibration lamp
- software needed to compare the performance of the instrument as determined by the model with actual calibration data.
- an optimizer tool which finds, for any given observation, the actual configuration of the instrument.

With this method we can close the loop between model description and actual observations resulting in a high fidelity dispersion solution for STIS high resolution echelle modes.

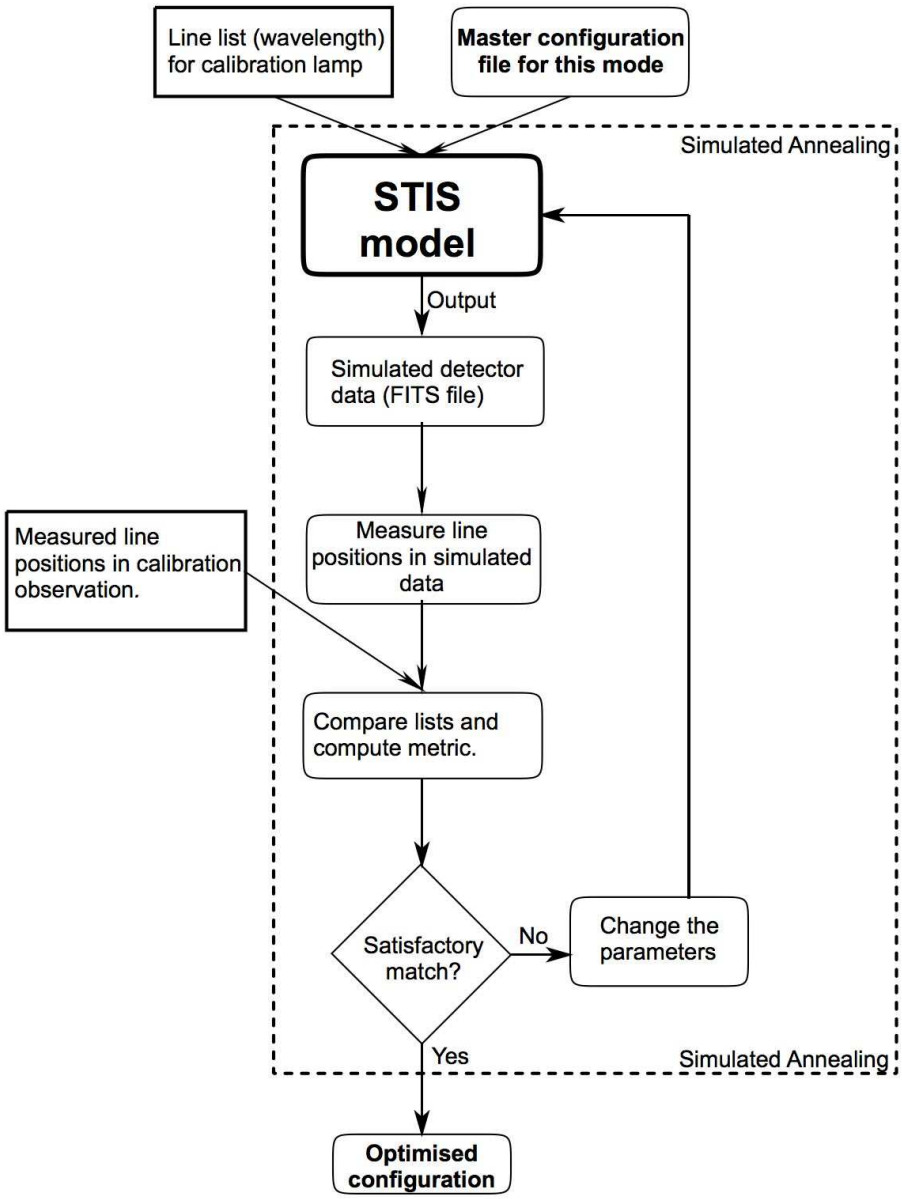

Figure 1: Schematic representation of the instrument model based approach used to derive a dispersion solution for the STIS high resolution echelle modes. An accurate list of wavelengths emitted by the Pt/Cr-Ne calibration source is a fundamental ingredient to the model. Starting from a master configuration for a given mode the predictions of the model are compared with the observed calibration data. The actual configuration of the instrument is then found in an iterative manner.

2.2. Input for the Instrument Model: The Lamp Project

Based on the experience with the International Ultraviolet Explorer (IUE) all spectrographs onboard HST carry a Pt(/Cr)-Ne hollow cathode lamp as calibration source. The HST has benefited very much from the work of the Atomic Spectroscopy Group at NIST. After a study at NIST had revealed that the best available wavelength values for Pt dated back to the 1930s and that their quality was not sufficient for use with Goddard High Resolution Spectrograph (GHRS), they took the initiative and measured from laboratory data Pt wavelength standards at the level of accuracy compatible with the spectral resolution (a few parts in 10^6) of the GHRS. The resulting line list and spectral atlas of a Pt-Ne lamp (Reader et al. 1990) contains accurate values for about 3000 lines.

The spectral output of Pt-Ne lamps is fully adequate for pure UV instruments like IUE or GHRS, with its wavelength range of 113 nm to 310 nm, where Pt has numerous emission lines. However, in order to cover the extended range of UV-optical spectrographs such as the Faint Object Spectrograph (FOS) and STIS, the addition of about 10 % Cr to the cathode is required. Pt/Cr-Ne lamps provide a continuous distribution of suitable emission lines for the range 115 nm to 800 nm. Although both FOS and STIS carried the Pt/Cr-Ne lamps, the pure Pt line list has been the basis for wavelength calibration of all spectrographs onboard HST.

For STIS-CE we required a complete list of the lines emitted by the Pt/Cr-Ne hollow cathode lamp in order to realize the full potential of the instrument model. Literature data for Cr proved inadequate for the purpose. Therefore, ST-ECF and NIST launched a project with the objective of obtaining wavelength standards from 115 nm to 320 nm in the spectrum of a Pt/Cr-Ne lamp. The main results of this project are line lists of the Pt/Cr-Ne spectrum (Sansonetti et al. 2004, Kerber et al., 2006 in preparation). The wavelengths in these lists are accurate to better than 1/1000 nm. More details of this work are also described in Kerber et al. (2004) and in two papers in this volume (Kerber et al. 2006a, 2006b).

2.3. Laboratory Data for the Virtual Observatory

We expect that these data will be valuable for applications beyond their original purpose, the wavelength calibration of STIS. In particular the accurate wavelength data on Cr will be useful for other fields such as the study of the interstellar medium (ISM) to which the HST and STIS have provided many important observations in the past, see e.g. Roth & Blades (1997). Two factors are important in this context: metallicity and depletion from the gas phase onto grains. Cr is a particularly relevant element since about 90 % of the Cr is missing from the gas phase. Accurate wavelengths for more than 5000 Cr lines will be a welcome addition for investigations of the ISM as well as for studies of stellar abundances. Another application for highly accurate wavelengths of Cr lines could be in fundamental physics. The constancy of the fine structure constant α has been under debate for quite some time and investigations from laboratory measurements and astrophysics have given conflicting results. Cr is one of the elements that has prominent lines in the absorption systems of quasars used for astrophysical analysis. Details of such potential applications are described in more detail in an upcoming paper by Kerber et al. (2006 in preparation).

ST-ECF and NIST are making all relevant results of the Pt/Cr-Ne lamp project available to the community (www.wavecal.org) also in a form compatible with the requirements of the Virtual Observatory (VO). To our knowledge this is a first in terms of such an extensive spectral data set obtained from laboratory measurements.

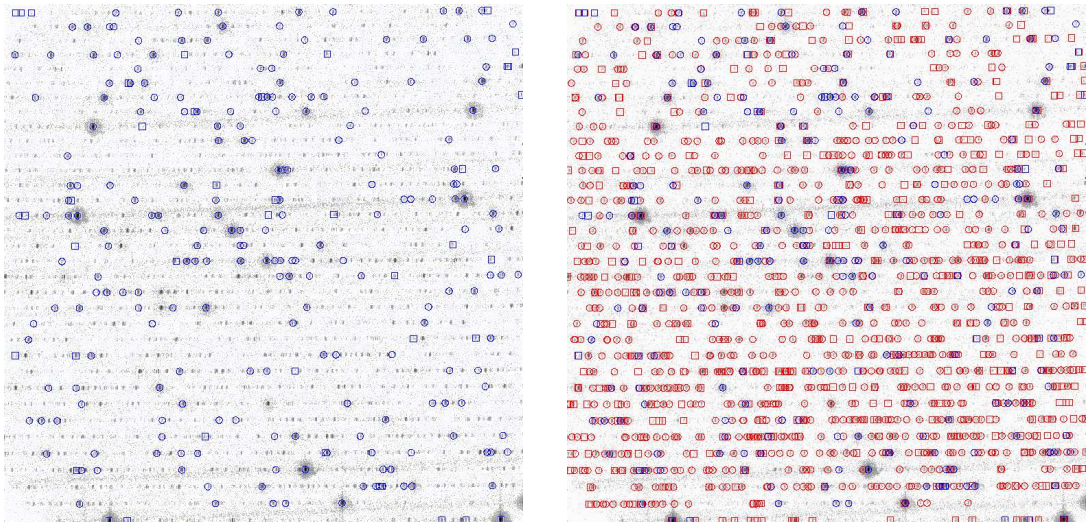

Figure 2: 2-dimensional STIS frame (echelle mode G230H) with a central wavelength of about 251.3 nm. Each order covers about 1.1 nm. The left panel shows the lines identified by the NIST Pt atlas (1990), the right panel illustrates the impact of the new Pt/Cr-lamp line list. In this particular frame the old list provided 258 lines in the wavelength range covered, whereas the new list has 1612 entries.

3. Application to STIS

STIS is the first HST spectrograph to provide both high resolution modes (R = 100 000) and large spectral coverage (115 nm to 1100 nm). For proper calibration of the STIS Echelle modes it is vital to use a large number of lines all across the 2-D detector array and to have a reasonable number of lines in every order. As a result of the ST-ECF/NIST Pt/Cr-Ne lamp project such information is now available. In particular we have added data for about 5500 Cr lines, which could not be used for wavelength calibration previously. A combination of the NIST Pt atlas and the new line lists provides about 11500 lines for calibration purposes in the region 111.5 nm to 433.2 nm. This forms the basis for the STIS-CE model-based wavelength calibration. The impact of the Cr lines is most pronounced in the near UV where up to 90 % of the lines observed are from Cr. Figure 2 gives an example showing a calibration exposure taken at a central wavelength of 251.3 nm. These data enable us to produce a dispersion solution that is more accurate and homogeneous across the whole STIS array than an empirical calibration.

4. Science Case

We carefully selected a science case that would both be relevant to typical uses of STIS and would demonstrate the accuracy of the model based solution. In this context it is important to note that the standard calibration pipeline for STIS provides very good science products based on careful calibration using empirical methods. Hence, CALSTIS calibration provides a very demanding benchmark and any improvement with respect to CALSTIS should be considered a significant achievement.
The specific case we selected makes use of data obtained for proposals 8043 and 8484 (PI: E.B. Jenkins). Their goal was to study the thermal pressure in the interstellar medium (ISM) using absorption lines from neutral carbon. They used a total of 21 lines of sight towards early-type stars. Observations were performed in both high resolution echelle modes E140H and E230H. They specifically aimed for the best possible wavelength resolution

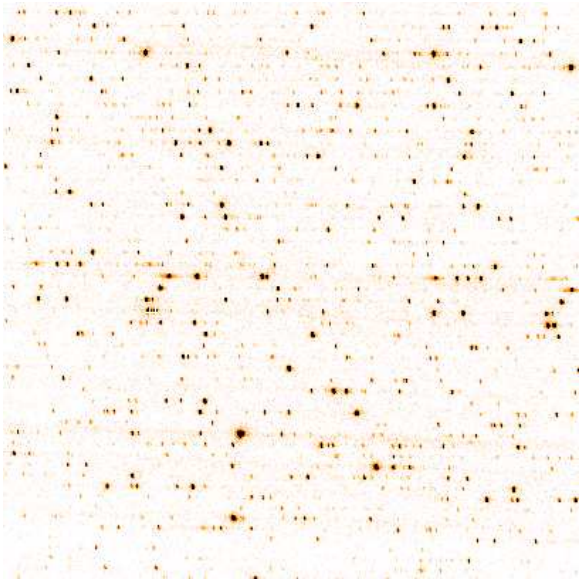

Figure 3: STIS 2-D array showing a deep wavelength calibration exposure (mode E140H, central wavelength 130.7 nm) taken with a Pt/Cr-Ne calibration lamp. On both sides of the array there are substantial overlap regions where the same line is apparent in two adjacent orders.

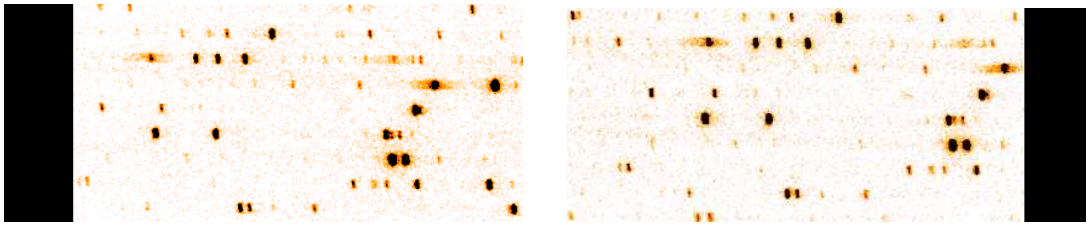

Figure 4: Enlarged section of the same exposure showing some overlap regions in detail. Several characteristic patterns of lines present in two adjacent orders can easily be discerned. For the science case we measured the difference in wavelength for interstellar absorption lines in these overlap regions.

and thus used the narrowest available slit (0.1 × 0.03 arcsecs) giving a $\lambda/\Delta\lambda$ of about 200,000 resulting in a FHWM of about 1.5 $\mathrm{km\,s^{-1}}$. The results of the program are described in Jenkins & Tripp (2001). The authors note the high quality of the CALSTIS pipeline reduction but also mention small offsets between the orders.

This kind of study is typical of the kind of science done with the high resolution spectrographs onboard HST. The fact that the observations were done at the highest resolution possible with STIS makes them especially suited for demonstrating the capabilities of any wavelength calibration procedure.

4.1. Overlap Regions in the STIS Spectral Format

In order to assess the quality of the model based dispersion solution we used interstellar absorption lines. The spectral format of STIS results in regions of considerable size at both sides of the 2-D array where orders overlap. A wavelength present in this overlap region will produce two line images, one each in two adjacent orders. A perfect wavelength calibration would of course result in exactly the same wavelength assigned to both images of the same line. For illustration purposes we show an example of a deep wavelength calibration

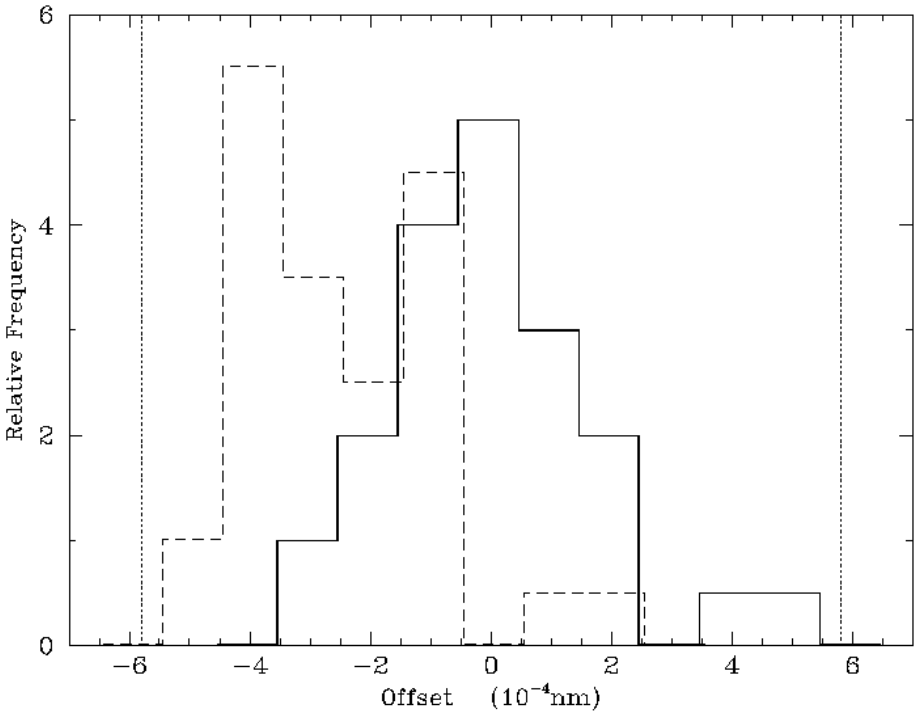

Figure 5: Histogram of the distribution of wavelength offsets measured between the same interstellar absorption line present in two adjacent order of the STIS array (mode E140H). The observations were made towards HD210839. The dotted vertical lines at $\pm$ 5.7 $\times$ 10^{-4} nm represent $\pm$ 1 pixel on the array. The dashed line shows the result for the data reduced with the CALSTIS pipeline. We find a systematic offset of about 0.5 pixel, which is within the specifications of the calibration. Data reduced with the STIS-CE dispersion solution (solid line) show essentially zero offset (about 0.1 pixel). The use of STIS-CE thus results in a factor of five improvement across the array.

exposure using the Pt/Cr-Ne lamp for mode E140H at a central wavelength of about 130.7 nm (Fig. 3). Enlarged versions of two overlap regions, one from the left and one from the right hand side of the array are shown in Fig. 4. Note the black regions beyond the limit of the array. Several characteristic patterns of lines present in two adjacent orders of the overlap region can easily be discerned. For the science case we measured the wavelengths of all available interstellar absorption lines in the overlap regions on the same frame once calibrated by CALSTIS and once calibrated using the STIS-CE model based dispersion solution. In order to make a fair comparison, all other steps of the data reduction remain the same. In total we analyzed three different lines of sights towards hot stars. In each frame we used about 50 individual measurements.

4.2. Results

We used the CALSTIS reduced frames as our reference and we found that in all cases studied the wavelengths of the same lines measured in adjacent orders show a systematic offset from zero. In the case of HD210839 (Fig. 5) we find a difference of about (-3.3 ± 1.9) $\times$ 10^{-4} nm or about 0.5 pixel for CALSTIS reduced data, which is within the stated accuracy of the wavelength calibration for E140H. For the STIS-CE dispersion solution we

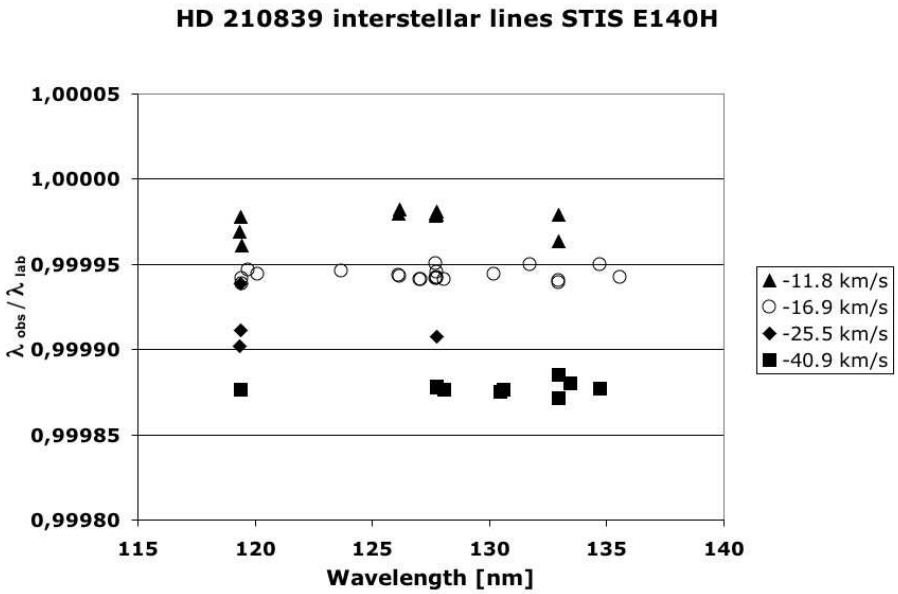

Figure 6: Velocities of interstellar absorption lines seen towards HD210839

find an offset of $(-0.6 \pm 1.7) \times 10^{-4}$ nm equivalent to 0.1 pixel. For the two other lines of sight studied (HD108 and HD203374A) the results are very similar, although the signal to noise ratio is less in both cases resulting in a slightly reduced accuracy. In all three cases a factor of five is gained in accuracy – over the calibration provided by the standard CALSTIS pipeline – by using the STIS-CE model based dispersion solution.

The scatter of the distribution remains essentially the same since the data were reduced in an identical manner expect for the dispersion solution. We also find a pattern in the offsets as a function of wavelength or position on the array. We suspect that the extraction of the spectrum which uses the empirical polynomial description contributes to this behavior. We intend to study this issue in more detail once the STIS-CE solution has been incorporated into the pipeline. In collaboration with the Spectrographs Group at STScI we are working towards an implementation of the STIS-CE calibration solution into the pipeline.

Using the new STIS-CE wavelength calibration we have looked at the interstellar absorption lines present in the E140H spectra of HD210839. Figure 6 shows the presence of four distinct velocities in the data. The features at velocity $-16.9\,\mathrm{km\,s^{-1}}$ are from neutral carbon and coincide with the strongest peak found by Jenkins & Tripp (2001). The other velocities result from other species in the ISM.

As mentioned before, STIS is arguably the best calibrated spectrograph of its kind. The fact that for STIS-CE we have been able to achieve a factor of five improvement in wavelength accuracy for the high resolution echelle mode E140H by using a dispersion solution based on a physical model of STIS, clearly demonstrates the excellent performance of instrument modeling techniques.

We hope that for future space and ground based instrumentation these methods will be widely employed in order to deliver the best possible calibration to the scientific user community.

5. Summary

- In the context of the ST-ECF's STIS calibration enhancement effort we have derived dispersion solutions for the high resolution echelle modes of STIS using a physical model of the instrument.
- In collaboration with the Atomic Spectroscopy Group at NIST we have measured the spectrum of Pt/Cr-Ne lamps similar to the calibration sources on STIS. The resulting line lists serve as high quality input for the STIS-CE model based calibration.
- All relevant results of the lamp project will be made available to the community also in a format compliant with the Virtual Observatory. In particular the new data for Cr is expected to be valuable for studies of the interstellar medium and also other fields of physics.
- Using a science case based on interstellar absorption lines we have demonstrated the gain in wavelength accuracy that can be achieved by using the STIS-CE physical model. For lines in the overlap regions of adjacent orders (E140H) we find a factor of five improvement over the already very good calibration provided by the standard CALSTIS pipeline.
- In collaboration with the STScI we intend to make the improved wavelength calibration for the STIS high resolution echelle modes available to the community and apply it to the STIS archive.
- The demonstrated success of the STIS physical model confirms the potential of this approach for superior calibration of scientific instruments both in space and on the ground.

Acknowledgments. This work was funded by the European Space Agency (ESA). We are grateful to D. Macchetto for his support throughout the project. We are also grateful to G. Nave, J. Reader, C. Sansonetti (NIST) for the excellent collaboration. We also thank the colleagues at STScI for their dedicated work on STIS and the good collaboration on the STIS-CE project. A very special "Thank you" goes to Anastasia Alexov and Mauro Fiorentino, former members of our group, for their contribution to the success of STIS Calibration Enhancement.

References

Ballester, P., Rosa, M.R., 1991, A&AS, 126, 563

Bristow, P., Kerber, F., Rosa, M.R., 2006, this volume, 299

Jenkins, E.B., Tripp, T.M., 2001, ApJS, 137, 297

Kerber, F., Lindler, D., Bristow, P., Lembke, D., Nave, G., Reader, J., Sansonetti, C.J., Heap, S.R., Rosa, M.R., Wood, H.J., 2006a, this volume, 324

Kerber, F., Bristow, P., Rosa, M.R., Nave, G., Reader, J., Sansonetti, C.J., and Lercher, G. 2006b, this volume, 318

Kerber, F., Rosa, M.R., Sansonetti, C.J., Reader, J., Nave, G., Bristow, P., Fiorentino, M., Lercher, G., 2004, SPIE 5488, 679

Reader, J., Acquista, N., Sansonetti, C.J., Sansonetti, J.E., 1990, ApJS, 72, 831

Rosa, M.R., 1995, in "ESO/ST-ECF Workshop on Calibrating and Understanding HST and ESO Instruments." P. Benvenuti (ed.), ESO Proceedings, #53, p.43

Roth, K.C., Blades, J.C., 1997, ApJ, 474, L 95

Sansonetti, C.J., Kerber, F., Reader, J., Rosa, M.R., 2004, ApJS, 153, 555

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Characterization of Pt/Cr-Ne Hollow-Cathode Lamps for Wavelength Standards in Space Astronomy

Florian Kerber, Paul Bristow, Michael R. Rosa[1]

Space Telescope European Co-ordinating Facility, Garching, Germany

Gillian Nave, Joseph Reader, Craig J. Sansonetti

National Institute of Standards and Technology, Gaithersburg, MD, USA

Georg Lercher

Sternwarte Max Valier, Karneid, Italy

Abstract.

The Space Telescope European Co-ordinating Facility (ST-ECF) and National Institute of Standards and Technology (NIST) are collaborating to study hollow cathode calibration lamps as used onboard the Hubble Space Telescope (HST). The project has two main components: First, we have observed the spectra of Pt/Cr-Ne lamps to obtain wavelengths for all emission lines between 115 nm and 320 nm. The observations were made at NIST using the 10.7-m normal-incidence spectrograph and a Fourier transform spectrometer (FTS) optimized for the vacuum ultraviolet (UV). The spectral region corresponds to the Space Telescope Imaging Spectrograph (STIS) Echelle modes. Wavelength calibration of all HST spectrographs has been based on the line list produced by Reader et al. (1990) using a Pt-Ne lamp, despite the fact that STIS and the Faint Object Spectrograph (FOS) use a Pt/Cr-Ne lamp. The addition of Cr is especially significant in the near UV where up to 90% of the observed lines are Cr. However, published Cr wavelengths are not sufficiently accurate for the calibration of STIS and FOS. Our work has established accurate wavelengths for more than 5000 Cr lines. The uncertainty is 2×10^{-4} nm for lines in the 110 nm to 180 nm region measured with the grating spectrograph and 10^{-5} nm for the 180 nm to 320 nm region observed with the FTS.

Second, we are attempting to better understand the performance of hollow cathode lamps and the physical processes involved in their long-term operation. Among the issues we have studied are the dependence of the spectrum on lamp current and cumulative operating time. We have performed accelerated aging tests that simulate operations on STIS using newly made space qualified lamps.

Our new Pt/Cr-Ne line list has been successfully applied in the STIS Calibration Enhancement effort; hence these results lead directly to an improvement in the quality of scientific data obtained from existing HST spectrographs.

We have also documented the sensitivity of the lamps' performance to optical alignment in the case of lamps that use a MgF_2 lens as window. Our findings constitute important lessons for the design and operation of future UV and optical spectrographs in space.

[1]affiliated to the Space Telescope Operations Division of Research and Scientific Support Department of the European Space Agency

1. Introduction

The selection of a light source for wavelength calibration of space observatories is based on practical considerations such as size, weight, power consumption, ruggedness and lifetime. The most important property of such a calibration lamp is that it has to provide a very rich spectrum at the resolution and in the range of the spectrograph. Based on these considerations a Pt(/Cr)-Ne hollow cathode lamp was chosen as the onboard emission line source for wavelength calibration of the spectrographs onboard HST: the Goddard High Resolution Spectrograph (GHRS), the Faint Object Spectrograph (FOS), the Space Telescope Imaging Spectrograph (STIS), and the Cosmic Origins Spectrograph (COS).

The high spectral resolution of the GHRS (a few parts in 10^6) demanded that wavelengths of the Pt emission lines be known to the same level of accuracy from laboratory data. A study at the National Institute of Standards and Technology (NIST) revealed that the best available values dated back to the 1930s and that their quality was not sufficient for use with GHRS. This fundamental issue was then addressed by Reader et al. (1990) who characterized the Pt-Ne lamp with accurate measurements of about 3000 lines.

The spectral output of Pt-Ne lamps is fully adequate for pure UV instruments like GHRS, with its wavelength range of 113 nm to 310 nm, where Pt has numerous emission lines. However, in order to cover the extended range of UV-optical spectrographs such as FOS and STIS, the addition of about 10 % Cr to the cathode is required. Pt/Cr-Ne lamps provide a continuous distribution of suitable emission lines for the range 115 nm to 800 nm. Although both FOS and STIS carried the Pt/Cr-Ne lamps, the pure Pt line list has been the basis for wavelength calibration of all spectrographs onboard HST.

STIS is the first HST spectrograph to provide both high resolution modes (R = 100 000) and large spectral coverage (115 nm to 1100 nm). For the STIS Calibration Enhancement (STIS-CE) effort, conducted at the ST-ECF, we use a physical instrument model which describes the configuration of the optical elements and the geometrical distortion; see Ballester & Rosa (1997) for details of this approach. In order to make full use of the capabilities of the model description for wavelength calibration, it is essential to employ a highly accurate and complete line list. For proper calibration of the STIS Echelle modes it is vital to use a large number of lines all across the 2-D detector and to have a reasonable number of lines in every order. This turned out to be impossible using the Pt-Ne line list since in some regions of the near UV the number of Pt or Ne lines is small and 90 % of the lines emitted by the lamp are from Cr. The existing literature data for Cr I (Kiess, 1953) and Cr II (Kiess, 1951) are not usable for STIS calibration enhancement purposes, since the Cr wavelengths are usually reported to only 0.001 nm and cover only the wavelength region redward of 200 nm.

It was therefore necessary to initiate a project with the objective of obtaining wavelength standards from 115 nm to 320 nm in the spectrum of a Pt/Cr-Ne lamp. Only then will we be able to achieve the wavelength scale accuracy STIS is potentially capable of. Additionally, and as a secondary objective, we aimed at gaining improved insight into the physical processes governing the performance and aging of such lamps; see also Kerber et al. (2006). Such knowledge is highly relevant for the continued operation of calibration lamps onboard HST, e.g. COS.

2. Laboratory Work

All of the laboratory work was done at NIST. In order to cover the wavelength range of the STIS Echelle modes - 115 nm to 320 nm - we used two instruments. For the far UV work we operated the 10.7-m normal-incidence vacuum spectrograph which records the spectrum on UV sensitive photographic plates (111.5 nm to 182.7 nm). This is the same instrument as the one used by Reader et al. (1990) to measure the wavelengths and intensities of the emission

lines of the Pt-Ne hollow cathode lamp. We have produced a line list containing more than 1200 lines; the line list and all other relevant details of the work are presented in Sansonetti et al. (2004). From the residuals in the fit of the wavelength interpolation polynomial to the standard lines, we estimate that the measured wavelengths have an uncertainty of about 2×10^{-4} nm.

For the near UV work we used the Fourier Transform Spectrometer (FTS) located at NIST's Synchrotron UV Radiation Facility (SURF). We recorded the lamp spectrum from about 155 nm to about 350 nm using two different photomultiplier detectors and a filter. Actually, the recorded spectrum extends beyond 400 nm, albeit at rapidly decreasing sensitivity. For the FTS spectra we achieved an accuracy of about 10^{-5} nm. For many Pt lines the wavelength accuracy is limited by partially or fully resolved hyperfine structure and isotope shifts. At the resolution of STIS, isotope splitting is not an issue, and the hyperfine structure can easily be taken into consideration for calibration purposes.

Exposure times at the 10.7-m spectrograph were normally 3 h at 20 mA and 24 h at 10 mA. These times produced approximately equal exposures for lines of Pt II. At the FTS we obtained between a few dozen and 2000 co-added spectra equivalent to exposure times of 30 min to 24 h.

3. Results and Application to STIS

Table 1 summarizes the results of both the far and near UV measurements. The new lists provide highly accurate wavelengths of the emission lines in the spectrum of a Pt/Cr-Ne lamp. In particular we have added data for about 5500 Cr lines, which could not be used for wavelength calibration previously. A combination of the NIST Pt atlas and the new line lists provides about 11500 lines for calibration purposes in the region 111.5 nm to 433.2 nm. This forms the basis for the STIS-CE model-based wavelength calibration.

Table 1: Summary of the line lists from the far UV and near UV part of the lamp project. More than 5000 Cr lines have been measured, identified, and are being used by STIS-CE.

	far UV (113.0 to 182.7) nm	FTS (155.0 to 433.2) nm
# of Pt lines	744	1280
# of Ne lines	105	590
# of Cr lines	258	≈5400
# of unidentified lines	128	<100
total # of lines	1235	7300

4. Dependence of Spectral Output upon Current

In our photographic spectra of the far UV region we recognized that the relative intensity of the neon lines with respect to the metal lines was significantly enhanced at lower lamp currents. In the near UV region the linear intensity response of the FTS enabled a more quantitative investigation of the change in the spectrum as a function of operating current. We took spectra at 7, 10 (standard current on STIS), 12, 15, 20, 25 and 30 mA.

Analysis of the spectra taken at different currents revealed a very well defined behavior of the line intensities as a function of current. For every spectrum, we calculated the ratio of the intensity of each observed line to the intensity of the same line in the 20 mA spectrum.

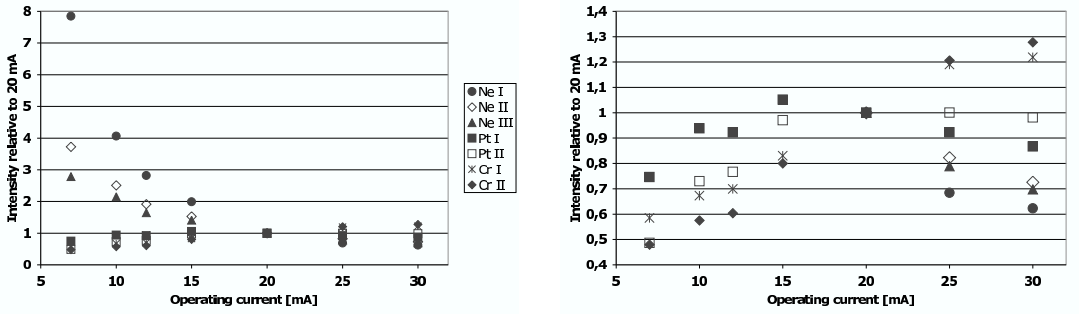

Figure 1: Variation of the line intensities as a function of the operating current of the lamp. Spectra were taken at 7, 10, 12, 15, 20, 25 and 30 mA. Note the pronounced difference in the behavior of the gas and metal lines (left panel). The curves of the two metals Pt and Cr show a less pronounced but still clear and characteristic difference in their response to varying current (right panel).

We then averaged the ratios for all lines of a particular species at each current. The results, which represent the average behavior of the line intensities as a function of current, are displayed in Fig. 1. There is a very pronounced distinction between the behavior of the metals and the gas lines. The gas lines show much higher intensities at lower currents. The metal lines (right panel) show a much more limited variation with current but careful analysis shows that the intensities of Pt and Cr lines vary in a distinct manner. A qualitative explanation of this behavior is given in Kerber et al. (2004).

5. Accelerated Aging Tests

An important secondary objective of the lamp project was to investigate the aging of Pt/Cr-Ne hollow cathode lamps and any changes in spectral output or operational characteristics associated with it. On STIS there is a total of three Pt/Cr-Ne lamps. During the 5.5 years of operation between installation in Feb 1997 and mid-2002, the LINE lamp had been used for a total of 117.5 h, accumulated during more than 8600 exposures (J. Valenti 2002, private communication). On average this corresponds to about 3.5 min of lamp usage per day on orbit. Most of the exposures taken are short; 91 % lasted less than 60 s, with an average exposure time of about 31 s. To simulate the working conditions of the STIS lamp we operated our lamps at 10 mA on an interval timer with 30 s on and 30 s off for several months. At several times during this period, and at the end of the aging period, we photographed the spectrum of the lamp on the normal-incidence spectrograph and compared it to a spectrum taken at the beginning of the test. Little change in line ratios or even total intensity has been found by us. A total of four Pt/Cr-Ne lamps have completed their aging test. They lasted between 1000 h and more than 2500 h. Lamp #16 had a depleted getter when we received it (barium coating inside the lamp that helps maintain a clean environment by trapping impurities) and showed an unusually rapid voltage increase from the beginning of the test. All lamps failed in a similar manner showing an anomalous discharge covering the first mica spacer. We speculate that the metal deposited there could form a conducting layer on the spacer leading to failure of the lamp.

The only obvious change observed before failure is an increase in operating voltage (Fig. 2), which could indicate a change in the rare gas pressure or in the cathode dimensions. Since the operating voltage seems to be the best diagnosis of lamp aging, and since it can easily be monitored, we strongly suggest that this voltage be part of the telemetry of future instruments and that it be used for analysis of the lamp's performance.

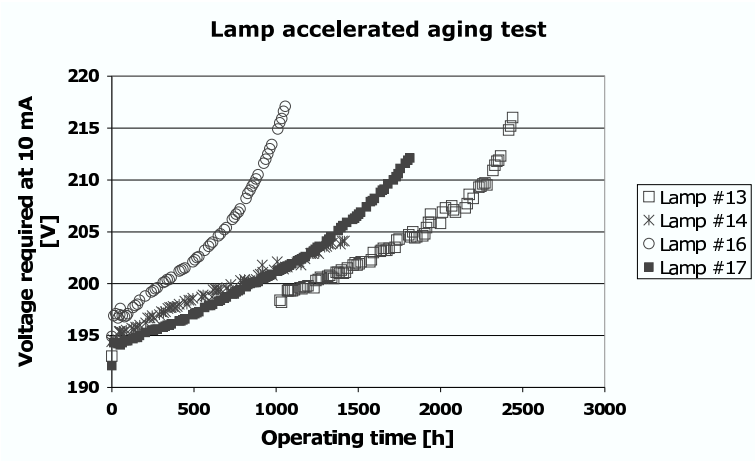

Figure 2: Change of the operating voltage - required for a current of 10 mA - as a function of accumulated operating time. The lamp was operated in a 30 s on / 30 s off duty cycle in order to mimic the operations onboard STIS. The lamps lasted for between 1000 h and 2500 h. All four lamps failed after experiencing an increase of 10 to 20 V in operating voltage.

6. Alignment for lamps with a MgF_2 lens

The lamps used on IUE, FOS, GHRS and STIS all have been ruggedized for use in space; this includes a special Kovar flange to mount the exit window, which consists of a single plano-convex lens made of MgF_2. Such a lens gives high transparency down to the cut-off of MgF_2 at 115 nm and focuses the light from the cathode. Since the index of refraction of MgF_2 is wavelength dependent and increases sharply below 150 nm, chromatic aberration is unavoidable with a single lens and perfect focus can be achieved at only a single wavelength.

For the laboratory measurements on the 10.7-m vacuum spectrograph, initial alignment was achieved visually by optimizing the illumination and brightness of the slit image as seen in the plate holder. With this alignment both the FOS and GHRS lamps produced reasonable spectra at long wavelengths. For the FOS lamp, however, the spectrum became progressively weaker toward shorter wavelengths, and below about 150 nm only the strongest lines were observed. Close inspection showed that the approximately focussed image of the hollow cathode discharge was moving off the side of the slit at short wavelengths. This suggests that the optical axis of the MgF_2 lens did not coincide with the axis of the discharge. The complete 5.5 h spectrum of the FOS lamp, at a highly compressed wavelength scale, is shown in the upper panel of Fig. 3. Quite evidently, careful alignment in the visible part of the spectrum was not sufficient to guarantee good alignment in the vacuum UV. A greatly improved spectrum was obtained, as shown in the lower panel of Figure 3, after we realigned the lamp using a photomultipler and the strong Pt II line at 124.8 nm. The intensity of the spectrum is increased by a factor 12 at 160 nm. At shorter wavelengths the improvement is even greater but cannot be accurately quantified because of the lack of comparison lines.

7. Summary

- We have measured the spectrum of Pt/Cr-Ne hollow cathode lamps in the range 115 nm - 433 nm adding accurate wavelengths for about 5000 Cr lines.

- This new line list has successfully been used as input for instrument model based dispersion solutions provided by the STIS Calibration Enhancement project.

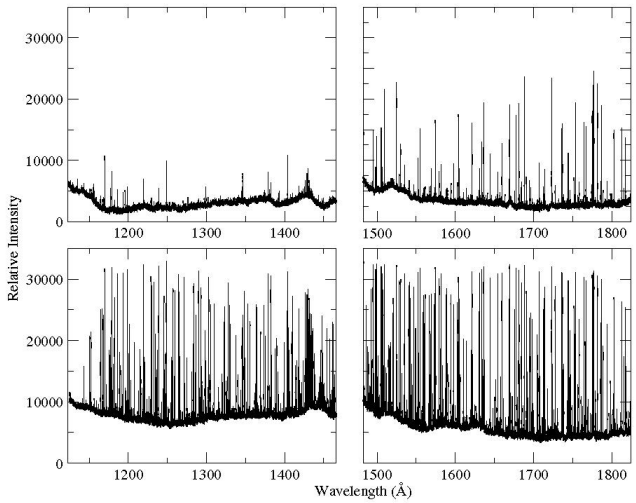

Figure 3: The spectrum of an FOS lamp after alignment in the visible (upper panel) and after alignment in the UV (lower panel). The dramatic increase in the line intensity at shorter wavelengths illustrates the strong sensitivity to alignment that results from use of a single element MgF_2 lens. See text for details.

- The lamps are very stable sources throughout their lifetime of about 1500 h at 10 mA during accelerated aging test that mimicked operations onboard STIS.
- The best indicator of lamp aging is an increasing operating voltage for a given current. We strongly recommend that this parameter be included in the standard telemetry for future instruments.
- MgF_2 singlet windows introduce chromatic alignment sensitivity that can lead to severe light loss if the lamps are not properly aligned. Alignment in the UV is essential.

Acknowledgments. This work was funded by the European Space Agency (ESA). We are grateful to D. Macchetto for his support throughout the project. We thank the HST project and the STIS IDT, in particular M. Niedner, B. Woodgate, T. Gull, R. Kimble for providing some of their back up lamps on loan.

References

Ballester, P., Rosa, M.R., 1991, A&AS, 126, 563

Kerber, F., Lindler, D., Bristow, P., Lembke, D., Nave, G., Reader, J., Sansonetti, C.J., Heap, S.R., Rosa, M.R., Wood, H.J., 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 324

Kerber, F., Rosa, M.R., Sansonetti, C.J., Reader, J., Nave, G., Bristow, P., Fiorentino, M., Lercher, G., 2004, SPIE 5488, 679

Kerber, F. & Wood, H.J., 2004, ST-ECF Newsletter, 35, 5

Kiess, C.C., 1953, J. Res. Natl. Bur. Stand. (US), 51, 247

Kiess, C.C., 1951, J. Res. Natl. Bur. Stand. (US), 47, 385

Reader, J., Acquista, N., Sansonetti, C.J., Sansonetti, J.E., 1990, ApJS, 72, 831

Sansonetti, C.J., Kerber, F., Reader, J., Rosa, M.R., 2004, ApJS, 153, 555

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Performance of the FOS and GHRS Pt/(Cr)-Ne Hollow-cathode Lamps after their Return from Space and Comparison with Archival Data

Florian Kerber

Space Telescope European Co-ordinating Facility, Garching, Germany

Don Lindler

Goddard Space Flight Center, NASA, Greenbelt, MD, USA

Paul Bristow, Dominik Lembke

Space Telescope European Co-ordinating Facility, Garching, Germany

Gillian Nave, Joseph Reader, Craig J. Sansonetti

National Institute of Standards and Technology, Gaithersburg, MD, USA

Sara R. Heap

Goddard Space Flight Center, NASA, Greenbelt, MD, USA

Michael R. Rosa[1]

Space Telescope European Co-ordinating Facility, Garching, Germany

H. John Wood

Goddard Space Flight Center, NASA, Greenbelt, MD, USA

Abstract. The Space Telescope European Co-ordinating Facility (ST-ECF) and National Institute of Standards and Technology (NIST) are collaborating to study hollow cathode calibration lamps as used onboard the Hubble Space Telescope (HST).

As part of the STIS Calibration Enhancement (STIS-CE) Project we are trying to improve our understanding of the performance of hollow cathode lamps and the physical processes involved in their long term operation. The original flight lamps from the Faint Object Spectrograph (FOS) and the Goddard High Resolution Spectrograph (GHRS) are the only lamps that have ever been returned to Earth after extended operation in space. We have taken spectra of all four lamps using NIST's 10.7-m normal-incidence spectrograph and Fourier transform spectrometer (FTS) optimized for use in the ultraviolet (UV). These spectra, together with spectra archived from six years of on-orbit operations and pre-launch spectra, provide a unique data set - covering a period of about 20 years - for studying aging effects in these lamps.

Our findings represent important lessons for the choice and design of calibration sources and their operation in future UV and optical spectrographs in space.

[1]affiliated to the Space Telescope Operations Division of Research and Scientific Support Department of the European Space Agency

1. Introduction

The two lamps flown on the GHRS and on the FOS are the only calibration lamps ever returned from space after extended use in orbit. All four lamps had been retired to exhibits at museums. In the case of the GHRS the instrument had been disassembled after return and the lamps were on display at the Fiske Planetarium in Boulder, Colorado. The FOS lamps were still inside the instrument which is on display at the National Air and Space Museum in Washington, DC. Courtesy of the HST project and both museums, we were able to obtain all lamps on loan. In the case of the FOS this required us to actually open up the FOS and extract the two lamps in a delicate operation (Kerber & Wood 2004).

The measurements we obtained yielded the only spectra of calibration lamps after return from space. Combined with pre-launch spectra recorded during ground testing in 1984 and the extensive collection of spectra obtained during the 6.5 years of orbital operation (Apr 1990 - Jan 1997) we have assembled a 20 year history of the spectral output of these lamps. This constitutes a unique database to study time dependent variations in the lamps.

2. Laboratory Work

All of the laboratory work was done at NIST. For the far UV work (Sansonetti et al. 2004) we used the 10.7-m normal-incidence vacuum spectrograph, which records the spectrum on UV sensitive photographic plates (111.5 to 182.7 nm). This is the same instrument used by Reader et al. (1990) to measure the wavelengths and intensities of the emission lines of the Pt-Ne hollow cathode lamp. For the near UV work (Kerber et al. 2006, in preparation) we used the Fourier Transform Spectrometer (FTS) located at NIST's Synchrotron UV Radiation Facility (SURF). We recorded the spectrum of the hollow cathode lamp from about 155 nm (instrumental cut-off) to about 400 nm using two different photomultiplier detectors and a filter. Although the recorded spectrum extends to 400 nm, sensitivity decreases rapidly beyond 350 nm.

Since we did not have any documentation on the wiring of the lamps, we had to establish the correct polarity by trial and error. After a few attempts all four lamps ignited and demonstrated their reliability 20 years after their first use and following a six-year period of hibernation in museums.

Exposure times on the 10.7-m spectrograph were 5.5 and 7 h for the GHRS and FOS lamps, respectively. On the FTS we coadded a few hundred spectra to obtain integration times of 4 to 6 h. The lamp current was strictly limited to 10 mA in order to match the operating conditions on the spacecraft and to minimize the risk of altering the properties of this unique hardware. In order to preserve the lifetime of the lamps, we did not attempt to obtain deep exposures. The resulting spectra contain about 700 lines for the GHRS Pt-Ne lamps, while the FOS Pt/Cr-Ne lamps yielded about 1500 lines.

Figure 1 shows an FTS spectrum and the corresponding section from the GHRS archive. A detailed analysis showed only limited variation in spectra taken on orbit and after return.

3. Analysis of the Spectra from Different Epochs

For the analysis we concentrate on the GHRS spectra because the resolutions of GHRS echelle modes and the FTS are similar, facilitating a direct comparison. We have combined spectra from the Instrument Definition Teams' pre-launch archive, the orbital archive, and the measurements at NIST, covering a time base of almost 20 years from 1984 to 2003.

First we looked for any changes in the intensity ratios of individual lines. To this end we carefully selected spectral regions which contain lines from both Ne and Pt, preferably also from different ionization stages. We chose regions at 185, 263 and 315 nm. By combining two or three echelle settings we used 15 - 25 lines in regions 5 - 10 nm wide. Figure 2

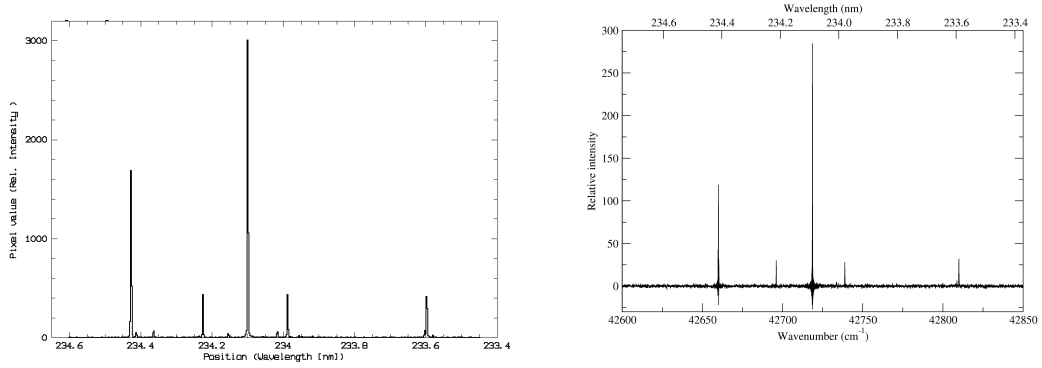

Figure 1: Sample spectrum of GHRS lamp #2.The wavelength range is approximately 233.4 – 234.6 nm. The spectrum in the left panel was taken with GHRS during orbital operations in 1994 (mode ECH B). The data were retrieved from the GHRS archive. The FTS spectrum (right panel) was obtained in the NIST laboratory in 2003.

illustrates a decrease of the observed intensity of the lines of about 25 % during the orbital phase. The lines belonging to a given element and ionization stage have been normalized to the intensity observed pre-launch and combined into an average for a given epoch. We used standard wavelength calibrations and so called SPYBAL's. A SPYBAL (SPectrum Y BALance) is a short calibration lamp exposure taken at standard wavelength settings to center the spectrum on the diode array in the cross-dispersion direction. A SPYBAL is normally taken whenever the optical element (e.g. grating) is changed. We selected wavelengths that had been observed as part of the pre-launch calibration program. The majority of the post-launch spectral calibration lamp observations are SPYBAL's. No significant differences between the individual species were found (Fig. 2). This indicates that the conditions in the discharge did not change noticeably over the years.

Comparisons between the three wavelength ranges selected show different rates of decrease at different wavelengths (Fig. 3). Based on the finding that no differences between elements and ionization stages exist, we have normalized all lines to the intensity observed pre-launch and combined into an average for a given epoch. The largest loss (approximately 45%) was found at the shortest wavelength of 185 nm. At 265 nm the decrease is about 25% while at 315 nm the behavior is not well described by a constant slope. While a decrease of the lamps' intensity over time would not be surprising, it is not evident that this should be wavelength dependent. Furthermore our accelerated aging tests (Kerber et al. 2004; 2006) on newly acquired but very similar lamps from the same manufacturer showed only very limited change in the lamps after more than 1000 h of operation. In contrast the GHRS lamps SC1 and SC2 only accumulated about 10 and 50 h respectively over a six-year period in orbit. To further elucidate the situation we made comparisons with other calibration exposures from the GHRS.

3.1. Flat Field Lamp

The GHRS used Xe lamps with major emission at 147 nm for a flat field illumination of the detectors. Light from the Xe lamps illuminated the detectors directly without any dispersing element in between. The flat field data show a monotonic decrease in intensity with almost constant slope (Fig. 4). Over the 6.5 years on orbit it amounted to a 15% loss at an effective wavelength of around 147 nm. While significant, this is far less than the 45% reported at 185 nm from the spectral lamps.

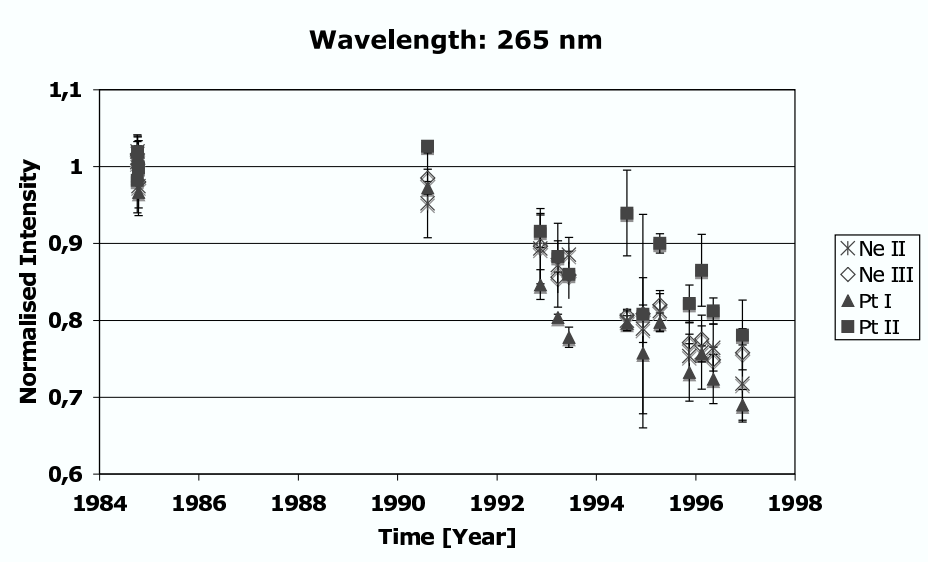

Figure 2: Change of the observed intensities of lines around 265 nm. Lines of a specific element and ionization stage have been normalized to the intensity observed pre-launch and combined into an average for a given epoch. A general decrease by about 25% is evident but no significant differences between individual species are apparent, indicating stable conditions in the discharge over the period investigated.

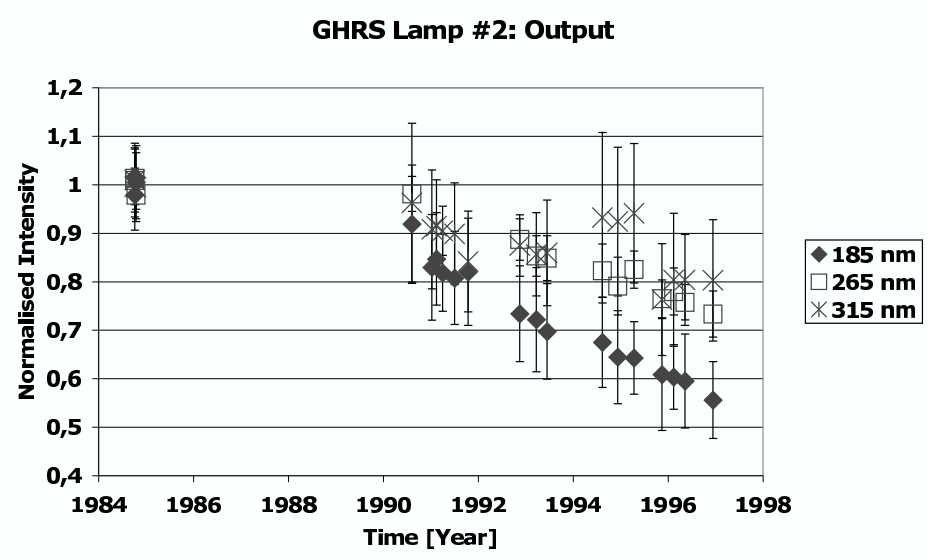

Figure 3: Change of the observed intensities of lines in three selected small (5-10 nm) wavelength bands distributed over the GHRS spectral range. Based on the finding that no differences between elements and ionization stages exist (Fig. 2) we have normalized all lines to the intensity observed pre-launch and combined into an average for a given epoch. The shortest wavelength shows a decrease of 45% in intensity while longer wavelengths show smaller losses.

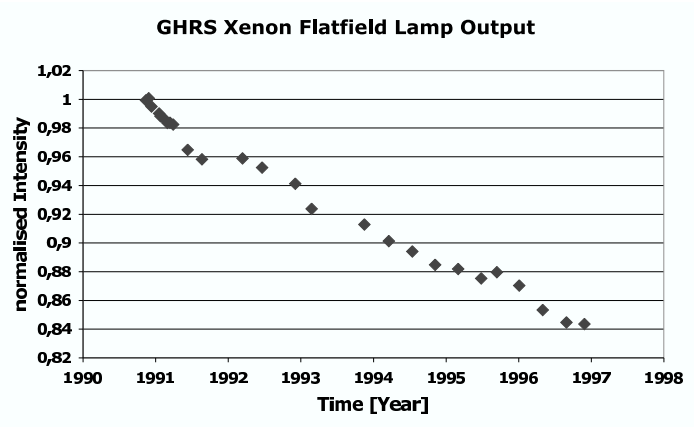

Figure 4: Change of the observed intensities of the GHRS Xe flat field lamp #2. Its main emission is around 147 nm. The intensity drops linearly over time, resulting in a 15% loss after 6.5 years.

3.2. External Calibration Source: Standard Star μ Columbae

For flux calibration, photometric standard stars were regularly observed. We have compiled a number of observations of μ Col over the complete orbital phase of GHRS. The flux of the star as observed by GHRS is constant except for a step observed in late 1993 when astronauts installed COSTAR during the first servicing mission to HST. COSTAR was an optical kit designed to compensate for the optical flaw of HST's primary mirror. Its components deliver the inverse of the spherical aberration, restoring the full imaging capabilities of the HST. COSTAR provided a highly improved point spread function (PSF) concentrating about 80% of the light in the central peak. This allowed more light to enter the slit of the HST's spectrographs while introducing two additional reflective surfaces in the optical path. In the case of the GHRS the net effect of COSTAR is illustrated by observations of the standard star μ Col (Fig. 5). The relative importance of the improved PSF (reduced slit losses) and added reflective losses explains the behavior after the end of 1993. At short wavelengths, the two additional reflections cause a net loss in sensitivity while at longer wavelengths the improved PSF leads to progressively better sensitivity. In Figure 5 three wavelengths closely corresponding to the wavelengths bands observed with the spectral lamps have been plotted. To facilitate comparison the fluxes in two of the bands have been multiplied by an appropriate factor. Except for the step, just explained, the flux of the standard star is highly constant.

4. GHRS Optical Layout and Interpretation of the Behavior of the Lamps

The observations of the standard star demonstrate that the performance of the elements of the light path for an external source were highly constant throughout the lifetime of the GHRS. The 15% decrease in intensity of the flat field lamps very likely originates in the lamps themselves since they illuminate the detectors directly without any optical elements in between (Fig. 6). The flat field lamps are Xe discharge lamps with a design not unlike the Pt-Ne spectral lamps. Light from the calibration lamps is reflected by two folding mirrors (Fig. 6, only one of the two mirrors shown) before it enters the spectrograph slit and from there on it shares the full optical path with an external target such as μ Col. Given the fact that the conditions in the discharge also do not seem to change over time (Fig. 2) we have to conclude that the observed wavelength dependent changes arise from changes in the

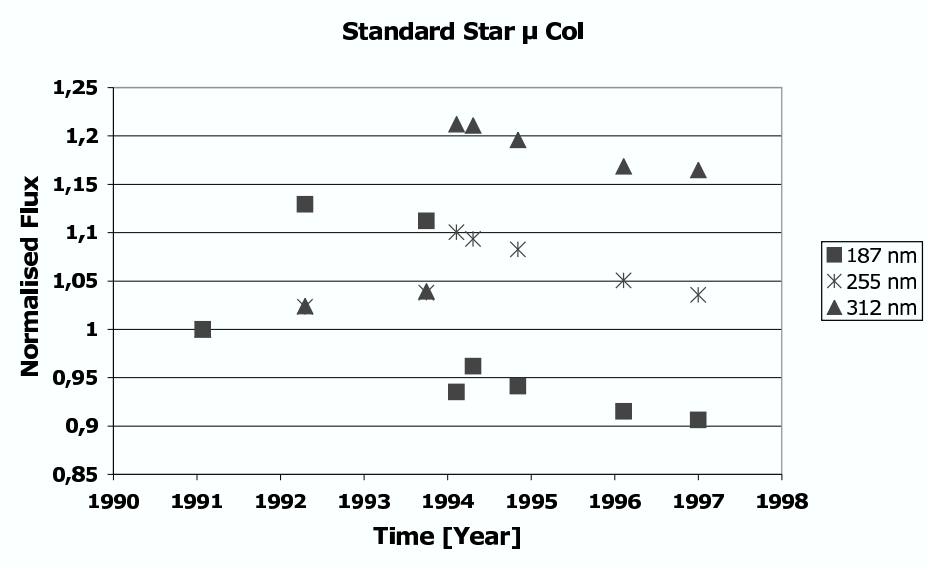

Figure 5: Observations of the standard star μ Col during GHRS orbital operations. Three wavelengths closely corresponding to the wavelengths bands observed with the spectral lamps (Fig. 3) have been plotted. To facilitate comparison the fluxes in two of the bands have been multiplied by an appropriate factor. Except for a step introduced by the addition of COSTAR (details in the text) the fluxes are constant to within 5%.

optical properties of either the lamp's window or the two MgF_2 surfaces directing the light to the spectrograph.

5. Summary

- We have assembled a unique data set spanning a period of 20 years consisting of spectra taken before launch, during operations in orbit, and after return of the four GHRS and FOS lamps.
- The lamps still work perfectly after 20 years and the conditions in the discharge seem to be constant in time, confirming that hollow cathode lamps are very reliable and stable sources for wavelength calibration.
- This stability is confirmed by accelerated aging tests of lamps similar to those used on GHRS, FOS and STIS extending to operating times well in excess of 1000 h.
- The data show some change in the intensities as observed over time. Careful analysis suggests that these can be attributed to optical components in the light path, specifically the two reflective prisms or the lamp windows.
- Hollow cathode lamps have previously been suggested as secondary radiometric standards (Klose, Hartig & Rosenberg, 1990) for monitoring changes in the optical path. Further investigations to this end are recommended.

Acknowledgments. It is a real pleasure to thank the HST project and in particular M. Niedner for support of this project and for making the lamps available to us. We would also like to thank D. Duncan from the Fiske Planetarium (Boulder, CO) and D. DeVorkin and his colleagues at the National Air and Space Museum (Washington, DC). Funding for this project is provided by the European Space Agency (ESA). We especially thank D. Macchetto for his support.

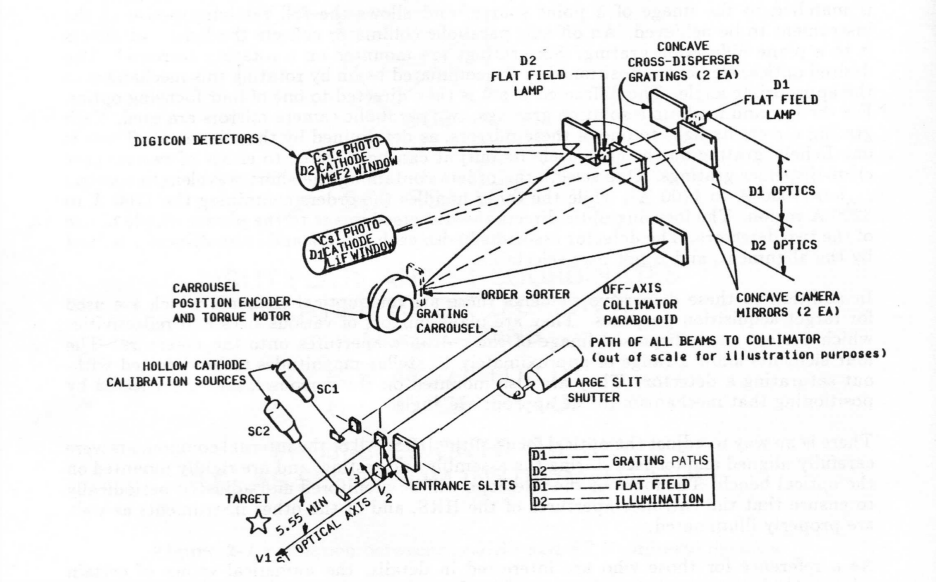

Figure 6: Schematic drawing of the GHRS optical system. (Not all optical elements shown.)

References

Kerber, F., Bristow, P., Rosa, M.R., Nave, G., Reader, J., Sansonetti, C.J., Lercher, G., 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 318

Kerber, F., Rosa, M.R., Sansonetti, C.J., Reader, J., Nave, G., Bristow, P., Fiorentino, M., Lercher, G., 2004, SPIE 5488, 679

Kerber, F. & Wood, H.J., 2004, ST-ECF Newsletter, 35, 5

Klose, J.Z., Hartig, G.F., Rosenberg, W.J., 1990, Applied Optics, 29, 2951

Reader, J., Acquista, N., Sansonetti, C.J., Sansonetti, J.E., 1990, ApJS, 72, 831

Sansonetti, C.J., Kerber, F., Reader, J., Rosa, M.R., 2004, ApJS, 153, 555

Part 4. Future Instruments: WFC3, COS and JWST

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Calibration Status and Results for Wide Field Camera 3

Randy A. Kimble

NASA, Goddard Space Flight Center, Greenbelt, MD 20771

and the WFC3 Team

Abstract. Wide Field Camera 3 (WFC3) is a general-purpose imager in development for installation in HST Servicing Mission 4 (SM4). Covering the wavelength range of 200-1700 nm in two observing channels, WFC3 offers powerful new capabilities, particularly in the near-ultraviolet and near-infrared bands. During 2004, the instrument was integrated and underwent a substantial suite of end-to-end characterization and performance tests. In this paper, we present a brief overview of the design and scientific purpose of WFC3, summarize the results of its test program to date, and highlight some recent developments in detector technology that will further enhance the performance of WFC3's IR channel.

1. Design and Purpose of WFC3

The WFC3 program originated in 1997 when NASA decided that the target date for the end of the HST mission would be moved from 2005 to 2010. Recognizing that this date would place HST's complement of imagers beyond their nominal design lifetimes, the HST Project initiated the development of a new camera to ensure and extend HST's imaging capabilities into those later years. WFC3 was thus conceived, to be developed as a facility instrument on behalf of the astronomical community (that is, with no PI team and no Guaranteed Time Observations). Day-to-day scientific guidance is provided by a Science Integrated Product Team from NASA/GSFC and STScI, with broad guidance exercised by a community-based Science Oversight Committee, chaired by R. O'Connell of the University of Virginia.

As problems with NASA's shuttle program have delayed the planned SM4 until at least late 2007 (with end-of-mission in ≈2013), the need for a backup camera for HST has, if anything, increased, but WFC3 provides much more than imaging "insurance" for Hubble. WFC3 is a powerful scientific imager in its own right, offering unique capabilities and scientific potential, particularly in the near-ultraviolet (NUV) and near-infrared (NIR) bands (very nicely complementary to the *I*-band-optimized Advanced Camera for Surveys).

WFC3 features two channels. The UV/Visible (UVIS) channel covers 200-1000 nm and is based on a two-CCD mosaic (4096×4096 total pixels). The CCDs are low-noise, UV-enhanced devices manufactured by e2v (formerly Marconi), yielding an unprecedented combination of sensitivity and field of view for NUV imaging with HST. The UVIS channel offers a 160×160 arcsec field with 0.039 arcsec pixels, with a rich assortment of broad, medium, and narrowband filters and a UV grism.

The IR channel (800-1700 nm) incorporates a custom HgCdTe focal plane array (FPA), developed by Rockwell Scientific Company, with 1014×1014 imaging pixels. The novel 1700 nm cutoff of this array permits low-dark-current operation at temperatures that can be reached with WFC3's passive radiator and thermo-electric cooling, and minimizes the impact of thermal background radiation from the relatively warm HST environment. The IR channel covers a field of 139×123 arcsec at ≈0.130 arcsec sampling, also with a sizable

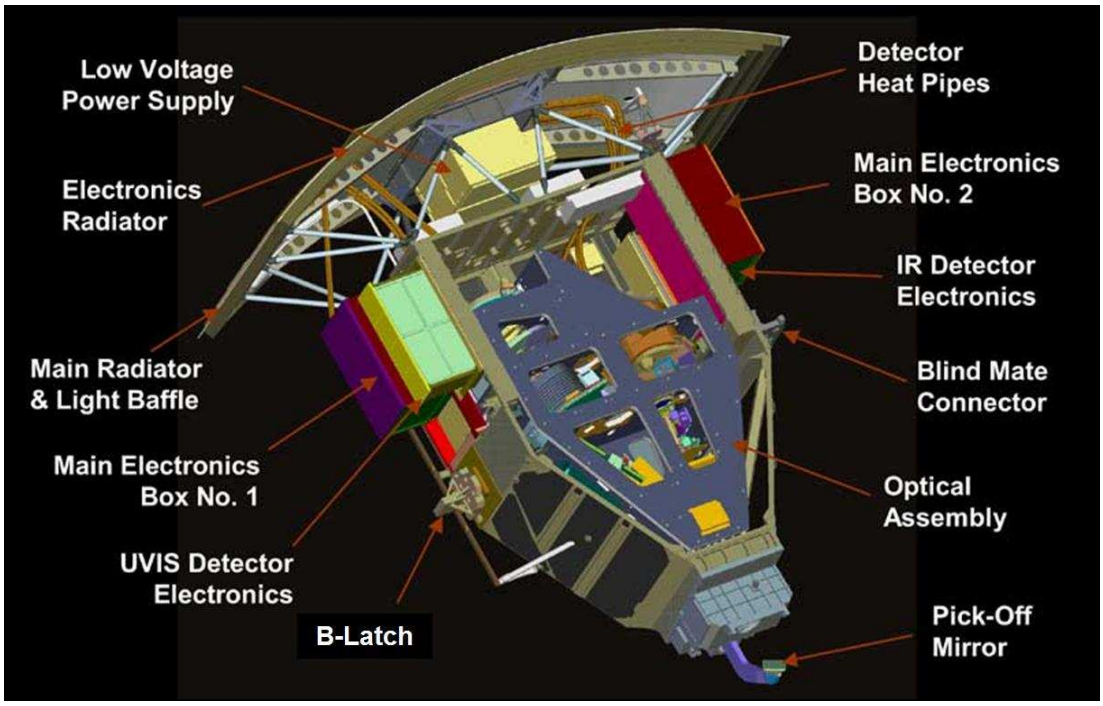

Figure 1: Schematic diagram of WFC3. Significant elements (radiator, enclosure, UVIS filter mechanism) are re-used WF/PC-1 hardware, while the electronics have substantial heritage from STIS, NICMOS, and ACS.

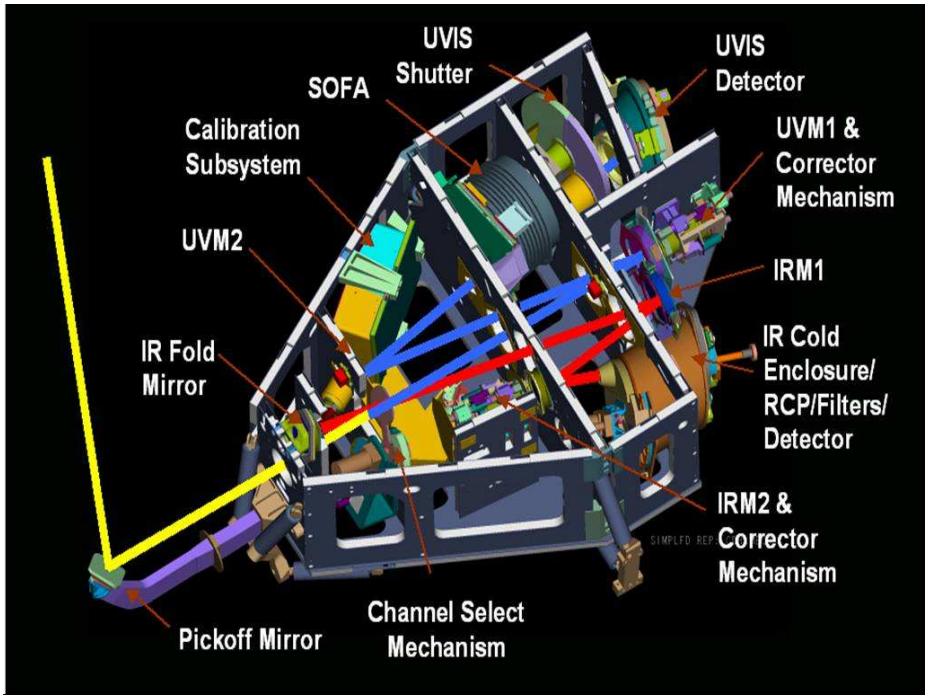

Figure 2: Optical layout of WFC3. The WFC3 pickoff mirror captures the on-axis beam from HST, which is then directed (as the shaded light beams in the figure indicate) to either the UVIS or IR channel, depending on the position of the Channel Select Mechanism.

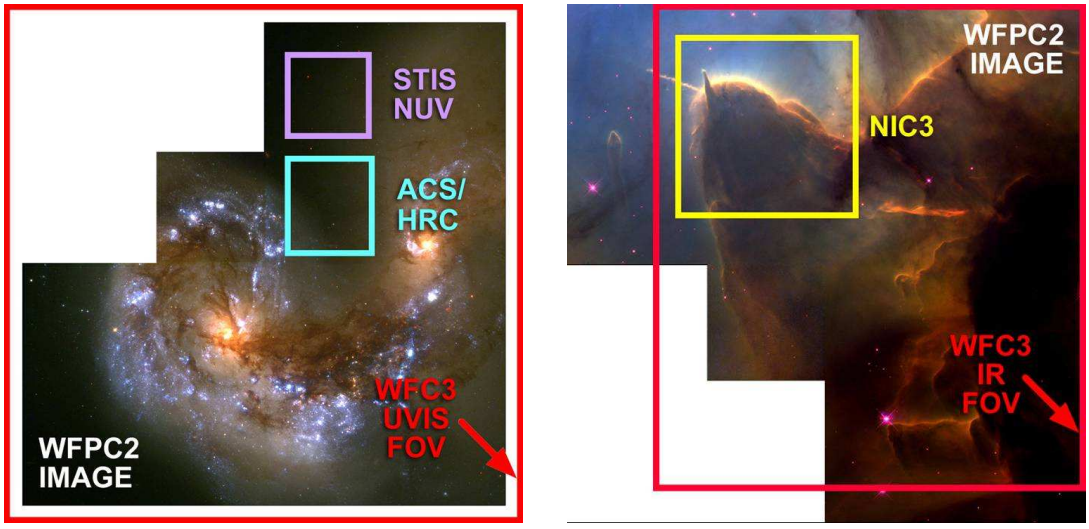

Figure 3: Schematic diagrams of the field of view (FOV) of the WFC3 UVIS (left) and IR (right) channels, compared with the FOV of current HST instruments. Note the large size of the UVIS FOV compared with those of the other HST imagers offering high NUV sensitivity (the STIS NUV and ACS HRC channels) and the dramatic increase in IR FOV vs. the largest field NICMOS camera.

complement of filters and grisms. The improved detector sensitivity and dramatic increase in field of view vs. NICMOS result in a major advance in NIR imaging for HST.

WFC3 is intended for installation into the radial instrument bay currently occupied by WFPC2 (see Figures 1 and 2 for schematic diagrams of the instrument's configuration). The dramatic advances in field of view (FOV) offered by WFC3 for NUV and NIR imaging are shown schematically in Figure 3. Coupling the FOV increases with the high sensitivity of the instrument at these wavelengths, WFC3 will enable a rich program of scientific observations. The UVIS channel will be particularly well suited to the study of the star formation history in nearby galaxies (by virtue of the high sensitivity of NUV observations to the age of stellar populations) and the chemical enrichment history of galaxies, and to the search for Ly α dropouts in the problematic redshift range of $z = 1 - 2$. The NUV portion of the UVIS bandpass will also probe one of the darkest spectral regions of the natural sky background, permitting the detection of very low surface brightness objects.

The IR channel will take advantage of the dark IR sky in space to study: Type Ia supernovae and the accelerating universe, galaxy formation at high redshifts, and the sources of cosmic re-ionization. The penetrating power of IR observations will be brought to bear on the dust-enshrouded processes of star formation, and the IR channel will also enable new studies of water and ices in our own solar system.

In short, WFC3 is a powerful new instrument, offering unique capabilities for astronomical observing. See Kimble, MacKenty, & O'Connell (2004) for additional details regarding the WFC3 design and scientific potential.

2. Highlights of Ambient and Thermal-Vacuum Testing

As is well known, there was a period throughout 2004 during which SM4 was officially cancelled on the grounds of shuttle safety. Currently, as directed by the NASA Administrator, the HST program is back to working toward a shuttle-based SM4, nominally scheduled for December 2007; a formal decision to execute the mission will not be made, however, until successful return-to-flight for the shuttle program.

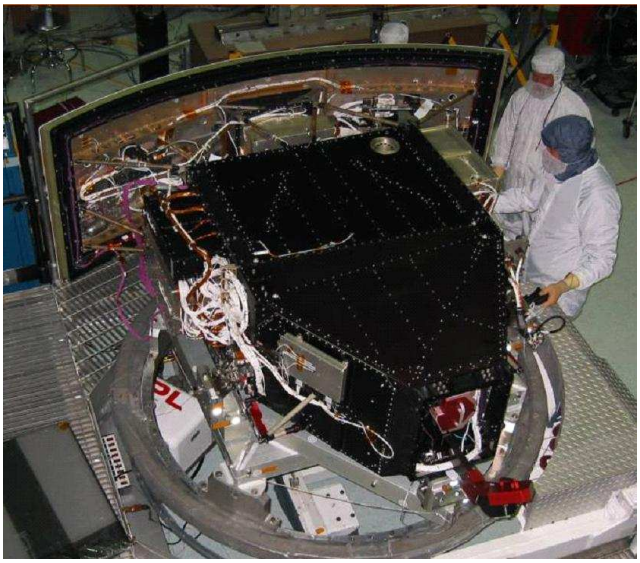

Figure 4: Photo of WFC3, integrated in preparation for its thermal-vacuum performance characterization test in late 2004.

Fortunately, the WFC3 program was not directed to stop work during 2004. Instead, the team was permitted to continue integration of the instrument (see Figure 4) toward execution of a "performance demonstration" thermal-vacuum test (with the idea that after demonstration of the capabilities of WFC3, it could potentially be made available for a non-HST launch opportunity). Though this meant that some issues that would normally be dealt with during instrument integration were deferred (and added to a "lien list" to be resolved now that the program has returned to a prepare-for-flight status), it enabled WFC3 to carry out an early and invaluable end-to-end evaluation of the instrument's performance.

The performance characterization thermal-vacuum test was carried out in August-October 2004 (with significant ambient testing of the UVIS channel performed before that time). Though the goals of the tests were not to produce a full scientific calibration, all aspects of WFC3 were investigated, with very satisfactory results overall. We highlight some of the key results here, in Tables 1 and 2 and Figures 5-8; additional details regarding the thermal-vacuum test program are presented by Bushouse et al. (2006) in this volume.

Table 1: Key UVIS Channel Thermal-Vacuum Results.

Characteristic	Specification; Goal	Measured
Dark current	<20 e^-/pix/hour	0.2-0.4 e^-/pix/hour
Read noise (rms)	<4 e^-/pix; <3 e^-/pix	2.98-3.08 e^-/pix
Linearity	<5% deviation over 100-50,000 e^-	<3% deviation
Encircled energy	250nm: >0.75; >0.80 in 0.″20 633nm: >0.75; >0.80 in 0.″25	250nm: 0.78-0.81 633nm: 0.77-0.81
Cal system	10,000 e^-/pix in <10 min Uniform to <2×	<1 min ≈7×
Filter ghosts	<0.2% of incident in a ghost	up to 15% for a small subset of filters (see text)

Table 2: Key IR Channel Thermal-Vacuum Results.

Characteristic	CEI spec; goal	Measured
Dark current	<0.4 e^-/pix/s; <0.1 e^-/pix/s	0.1 e^-/pix/s
Read noise (rms) (CDS pair of images)	<15 e^-/pix; <10 e^-/pix	≈23 e^-/pix
Full well	>100,000 e^-/pix; >150,000 e^-/pix	≈100,000 e^-/pix
Encircled energy (1000nm)	>0.56; >0.61 in 0.″25 >0.72; >0.80 in 0.″37	0.52-0.56[a] 0.73-0.77
Encircled energy (1600nm)	>0.48; >0.54 in 0.″25 >0.75; >0.80 in 0.″60	0.40-0.44[a] 0.77-0.81
Cal system	10,000 e^-/pix in <10 min Uniform to <2×	<1 min ≈25×
Filter ghosts	<0.2% of incident in a ghost	<0.2%

[a] The measured encircled energies require small corrections for the difference in central obscuration between the ground test aberrated beam simulator vs. the real HST Optical Telescope Assembly; with those corrections, the measured encircled energies are in line with the requirements (see Figure 5).

The detector performance in end-to-end thermal-vacuum testing summarized in Tables 1 and 2 is in complete agreement with the results of testing at the device and sub-system level: the UVIS CCDs show a low read noise of 3 e^- rms (the lowest of any HST CCD camera) and negligible dark current; the IR detectors were known to have higher-than-specified read noise, but the very low dark current vs. specifications and the measured throughput of the optical system bring the overall sensitivity of the channel within specification.

Optical performance is excellent (Figure 5), with many encircled energy results well exceeding specifications; slight departures seen in the IR channel are likely the result of a small amount of detector interpixel diffusion.

The throughput of the UVIS channel is in close agreement with pre-test predictions (based on individual component measurements) throughout the visible, and (delightfully) well exceeds those predictions in the NUV range, where our early CCD calibrations had underestimated the sensitivity of the UV-enhanced devices (the NUV test setup has since been improved). In the IR, the end-to-end throughput runs systematically 10-15% below component-based predictions across the band. This discrepancy is a bit beyond the expected error bars in the comparison, so the team is investigating to see whether there is a correctable source for the discrepancy. In any case, the measured end-to-end throughput in thermal-vacuum testing confirms the dramatic discovery potential for WFC3 in both the NUV and NIR bands; see Figure 8 for a tantalizing comparison of WFC3 vs. other HST instruments.

The most noteworthy performance discrepancies shown in the tables are the severe ghosting exhibited by a small subset of the UVIS filter set and poor performance of the internal calibration subsystem at this stage of integration. The filter ghosting issue is described in detail by Baggett et al. (2006) in this volume; to summarize briefly, we have procured greatly improved replacements for all high-priority filters with out-of-spec ghost intensities. The Scientific Oversight Committee has confirmed the acceptability of these replacements and the replacement process is in work. For the calibration subsystem, the UVIS channel simply requires improved alignment, while the IR channel flat-field non-uniformity will be rectified by replacing one of the transfer lenses in the calibration path.

Overall the thermal-vacuum characterization program was extremely successful and provides invaluable guidance in the team's ongoing preparation of the instrument for flight.

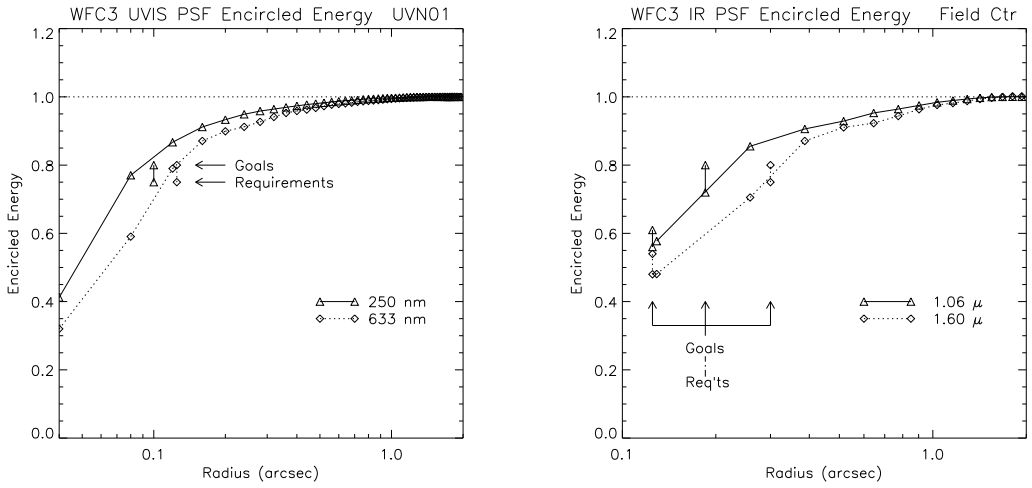

Figure 5: (left) Encircled energy vs. radius for UVIS channel images. (right) Encircled energy vs. radius for IR channel images. Phase-retrieval measurements of the delivered wavefront confirm the high quality of the optical system. Encircled energy results in the IR channel are likely affected by a small amount of interpixel diffusion in the detector.

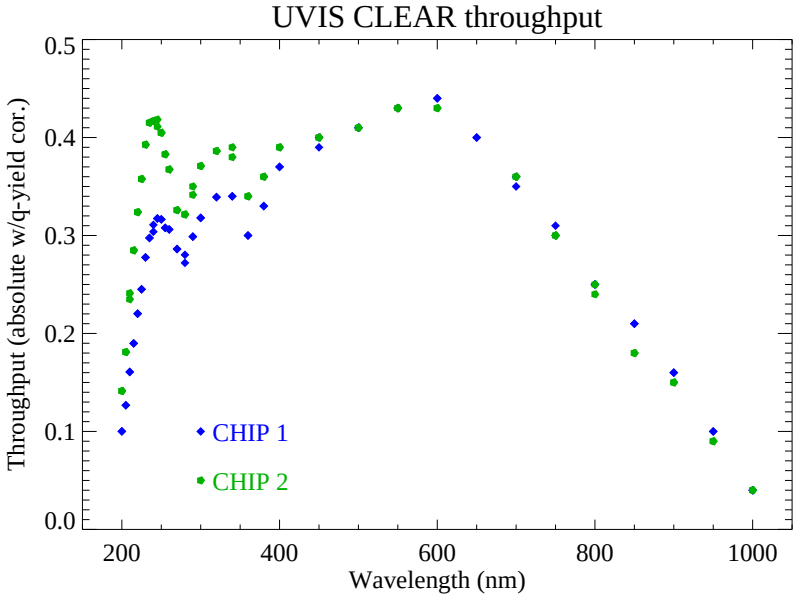

Figure 6: Unfiltered throughput (detected photons/input photons) for the UVIS channel; representative values are shown for each of the CCDs. The actual signal detected at short wavelengths (in electrons/input photon) is higher than shown here due to CCD quantum yield, which exceeds one electron per detected photon below ≈340nm. The curves shown here have been corrected for that effect.

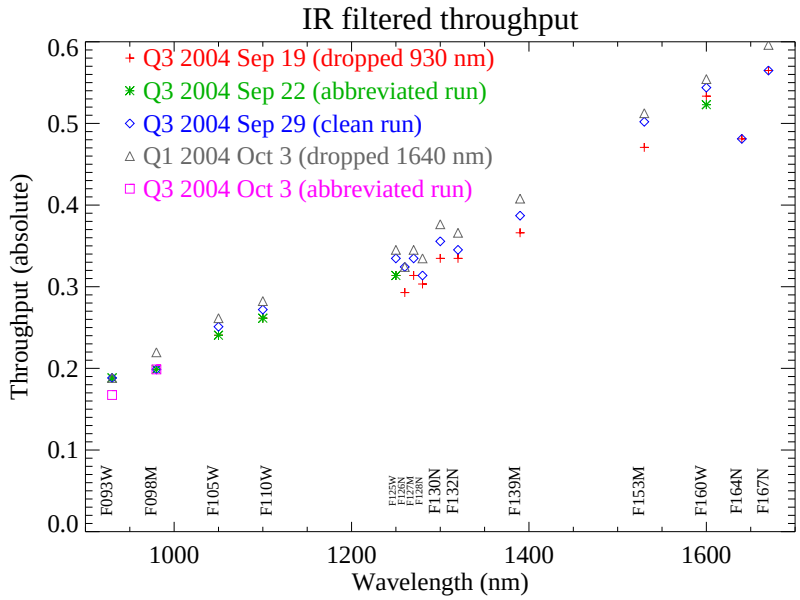

Figure 7: Throughput of IR channel as measured in thermal-vacuum test, including filters. The overall slope vs. wavelength is dominated by the detector quantum efficiency.

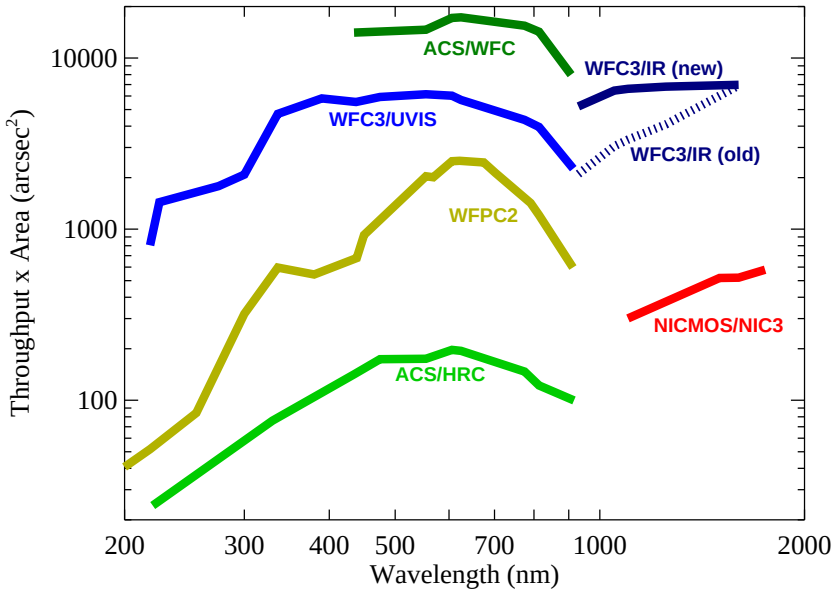

Figure 8: Comparative discovery efficiencies of various HST imagers. The discovery efficiency is defined as the throughput of the instrument (detected photons/input photons) $\times$ the area of the instrument's field of view. The curves are generated by calculating this quantity, including the filter transmission, at the center of band for the different broadband filters available for each instrument. For the WFC3/IR channel, the dashed curve (old) is based on the end-to-end throughput measured in the thermal-vacuum testing described above; the solid curve (new) scales up the throughput to the higher quantum efficiencies offered by new IR detectors that are currently in fabrication for flight (see Section 3). WFC3's uniquely powerful capabilities in the NUV and NIR bands are evident.

3. Improved Detectors for the IR Channel

In the course of development of WFC3, radiation testing of the IR detector arrays revealed an unexpected phenomenon. When exposed to a high-energy proton beam in the ground radiation test, the detector images showed a significant diffuse signal between the direct proton hits (Figure 9 - left). Further testing isolated the source of the diffuse signal to the thick CdZnTe substrate on which the HgCdTe detection layer is grown (Waczynski et al. 2005); when hit by high-energy particles, the substrate appears to produce luminescent output that is then detected by the HgCdTe layer.

Without a detailed understanding of the physical mechanism, scaling from the laboratory test conditions to the orbital radiation environment, and from one FPA to another, is uncertain. Our best estimate, however, is that the background produced in orbit by cosmic rays from this effect could be comparable to the intrinsic dark current of the devices, thus lowering the faint-target sensitivity of the IR channel if the effect is not eliminated.

Fortunately, in the time since our current flight FPA was acquired, Rockwell Scientific Company has successfully developed a process for removing the CdZnTe substrate after hybridizing the HgCdTe arrays. WFC3 radiation tests confirm that substrate removal completely eliminates the radiation-induced background phenomenon (Figure 9 - right). Hence, substrate-removed lots of WFC3 IR arrays are currently being procured, with the intention of packaging a new FPA for flight and removing the radiation risk.

Of even greater benefit, the substrate-removed FPAs show extremely high quantum efficiency (QE) across the WFC3 IR band, greatly improved over the performance of the current WFC3 flight device, FPA64 (Figure 10). (FPA64 was in the instrument for the thermal-vacuum testing described above.) The performance observed for the new devices appears to cluster in two groups: one with relatively flat QE across the WFC3 IR band, and one with more sloped QE, but both offering marked improvement over FPA64.

The new devices do have one drawback. So far, the substrate-removed devices produced for WFC3 exhibit higher dark current tails than seen in FPA64 (Figure 11). Interestingly, there is a clear correlation that the group of detectors with high, flat QE curves shows systematically larger dark current tails than the lower QE group.

Additional cooling of the detector produces marked dark current benefits. Our thermal-vacuum test demonstrated good thermal performance margin, so we can confidently project that the IR FPA can be operated at least down to 145K vs. the previous operating point of 150K. This 5 degree shift drops the dark current by more than a factor of two and significantly reduces the fraction of pixels in the dark current tail (Figure 12).

Fabrication of additional devices from the substrate-removed lots is ongoing, and Rockwell believes that small adjustments in the substrate-removal process for the remaining parts will yield better combinations of QE and dark current performance. However, we demonstrate below that even the currently available devices offer markedly improved scientific performance for WFC3 vs. FPA64, even neglecting that detector's radiation vulnerability.

In Table 3, we show a metric of "survey speed" in different filter bands, for several candidate WFC3 detectors (cf. Robberto et al. 2005). This metric is defined as the product of the speed (i.e. 1/exposure time required) to reach a given faint target magnitude in the given filter band $\times$ the area of the IR channel's field of view. The metric is then normalized so that the value for the NIC3 camera of NICMOS (the widest field imager of that instrument) is equal to 1.

The metric is tabulated for the J & H bands (F110W and F160W, respectively) and a representative narrow-band filter (F126N). The values for FPA64 demonstrate the order-of-magnitude increase in survey speed over NICMOS previously touted for the WFC3 IR channel (with our measured end-to-end throughputs). For each of the devices, we tabulate the mean dark rate for a particular percentage of the pixels, and we then scale the effective FOV for WFC3 by that percentage to account for the effects on surveying efficiency of masking the hot pixel tails.

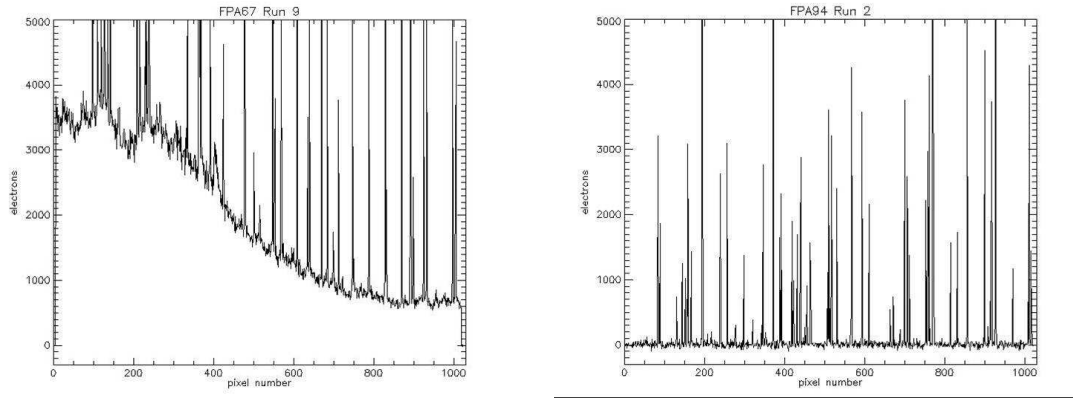

Figure 9: Radiation test results for WFC3 IR detectors. (left) cut across irradiated image for FPA with CdZnTe substrate, showing diffuse signal between the discrete hits; (right) analogous image cut for substrate-removed FPA, showing elimination of luminescence effect.

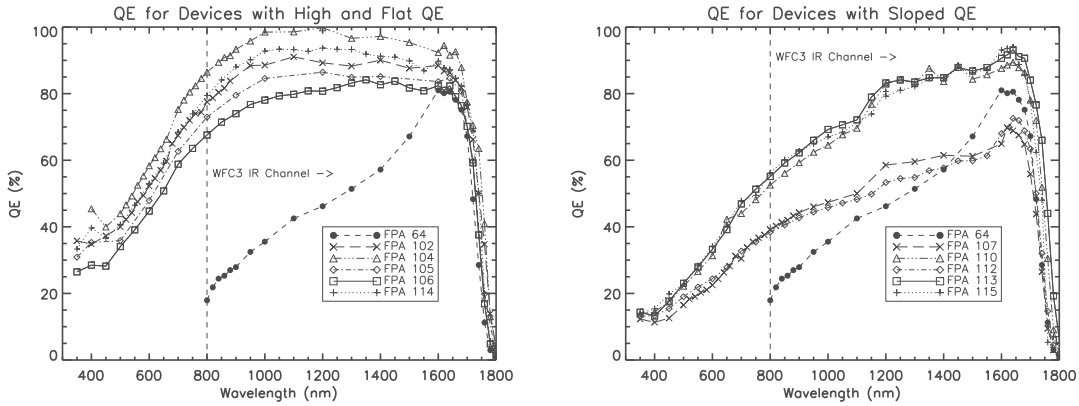

Figure 10: (left) Strikingly high QEs for group of substrate-removed IR FPAs, compared with QE of current flight detector, FPA 64. (right) QEs of other group of substrate-removed FPAs; less QE improvement, but better dark current performance than devices at left.

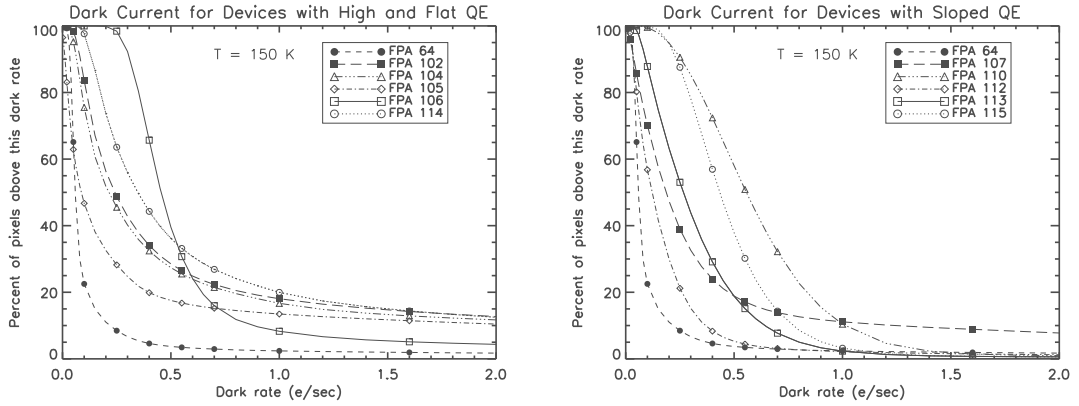

Figure 11: Cumulative dark current histograms for the two groups of detectors and current flight device, FPA64. At each value of dark current, the quantity plotted is the percentage of pixels with dark current exceeding that value. The group of substrate-removed devices with higher QE also exhibits a higher percentage of pixels in a high dark current tail.

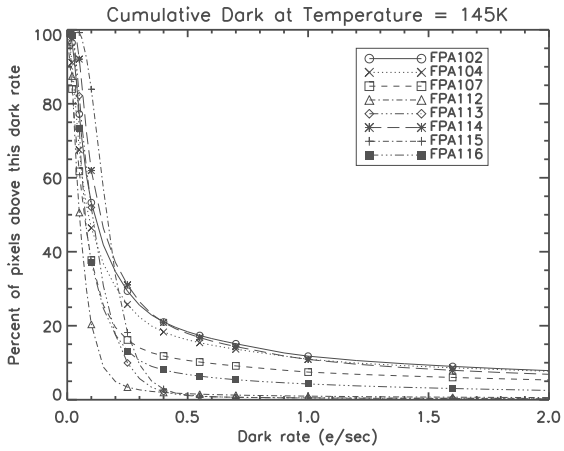

Figure 12: Cumulative Dark Current Histograms at 145K Operating Temperature

The results show that even a conservative choice of substrate-removed detector, such as FPA112 or FPA116 (with little or no hot pixel masking required), offers noticeable performance benefits over FPA64. The best high-QE parts (FPAs 104 and 105) offer more dramatically improved survey efficiency, though with greater complexity in observing (more dithering) and calibration to deal with the higher hot pixel percentages. Finally, the last line of the table shows the superb performance of a detector offering an "ideal" combination of parameters: 90% QE across the band along with low dark tail and read noise.

Table 3: Survey Speed Metric for Candidate IR Detectors.

Detector	CDS Noise	Mean Dark (145K)	F110W	F160W	F126N
FPA64	24	0.04 (100%)	12.5	9.1	11.7
FPA112	24	0.09 (100%)	15.6	7.7	12.4
FPA116	29	0.18 (98%)	20.0	7.4	13.0
FPA104	33	0.14 (90%)	28.7	10.4	18.2
FPA105	28	0.10 (90%)	26.6	10.0	20.6
Potential FPA?	20	0.04 (100%)	31.4	12.5	36.1

In conclusion, WFC3 will be a scientifically powerful addition to the HST suite of instruments. The WFC3 team looks forward to completing the instrument and providing this new capability to the astronomical community.

References

Baggett, S. et al., 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 348

Bushouse, H. et al.,2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 343

Kimble, R. A., MacKenty, J. W., & O'Connell, R. W. 2004, Proc. SPIE, 5487, 266

Robberto, M. et al. 2005 in Scientific Detectors for Astronomy, eds. P. Amico & J. W. Beletic (Dordrecht: Springer), in press

Waczynski, A. et al. 2005, Proc. SPIE, 5902, 161

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

WFC3 Calibration and Data Processing

H. Bushouse, S. Baggett, T. Brown, G. Hartig, B. Hilbert, J. Kim Quijano, J. MacKenty, I. N. Reid and M. Robberto

Space Telescope Science Institute

B. Hill, R. Kimble and O. Lupie

NASA/Goddard Space Flight Center

1. Introduction

Wide Field Camera 3 (WFC3), a panchromatic imager being developed for the Hubble Space Telescope (HST), is now fully integrated and over the past year has completed first rounds of extensive ground testing at Goddard Space Flight Center (GSFC), in both ambient and thermal-vacuum test environments. This report summarizes the results of those tests and describes the pipeline processing methods that will be used to calibrate WFC3 data.

WFC3 is designed to ensure that the superb imaging performance of HST is maintained through the end of the mission and takes advantage of recent developments in detector technology to provide new and unique capabilities for HST. WFC3 contains ultraviolet/visible (UVIS) and near-infrared (IR) imaging channels, offering high sensitivity and wide field of view over the broadest wavelength range of any HST instrument. It is slated to replace the current Wide Field and Planetary Camera 2 during Servicing Mission 4.

The WFC3 UVIS channel is based on elements from the Advanced Camera for Surveys (ACS) Wide Field Camera (WFC), with a 4096x4096 pixel Marconi CCD covering a 160x160 arcsecond field of view. The WFC3 UVIS channel is optimized for maximum sensitivity in the near-UV and contains a complement of 48 spectral filters and a grism. The WFC3 IR channel uses a 1024x1024 pixel HgCdTe Hawaii-1R detector array covering a 135x135 arcsecond field of view. The array sensitivity is optimized in the 0.8—1.7μm spectral range. The IR channel accomodates 15 filters and 2 grisms for slitless spectroscopy.

2. Test Setup and Procedures

WFC3 is designed to have its UVIS and IR detectors cooled to flight temperatures of -83°C and -123°C, respectively, to minimize dark current and thermal background. The IR detector can only be cooled sufficiently close to this flight temperature when it is in a thermal-vacuum environment, while the UVIS detector can be tested in both ambient and thermal-vacuum conditions. To date, the WFC3 UVIS channel has undergone two episodes of ambient testing, in January and June–July 2004. The UVIS and IR channels have together undergone thermal-vacuum testing during September–October 2004.

In both environments, the WFC3 is mounted to the Radial Instrument Alignment Facility (RIAF), which provides a reference to the HST latch plane. An optical stimulus apparatus, known as "CASTLE", is also mounted to the RIAF and is used to provide external sources to WFC3 for testing. The CASTLE is capable of providing point and extended targets, as well as flat-field illumination, over the entire wavelength range of WFC3. Single- and double-mode monochromators in CASTLE can be used to control the bandwidth of all sources. The CASTLE also includes NIST-calibrated detectors to measure absolute flux levels of incident sources, which are used to measure the absolute throughput and sensitivity of WFC3.

During routine data taking episodes, WFC3 test exposures are commanded via an HST-style SMS that is run on the instrument. Complementary scripts are used to control the CASTLE optical stimulus, to provide the desired source for each exposure, resulting in a highly automated process. All exposures are processed, previewed, and archived locally, and are also automatically sent to the STScI pipeline (OPUS) system to be converted to FITS files and stored in the long-term HST archive.

3. Ground Tests Performed

The goals of the ground tests performed to date were to:

- characterize the thermal performance of WFC3
- demonstrate flight-like operations of the UVIS and IR channels
- verify and characterize the science capabilities of WFC3

The types of tests that have been performed during the ambient and thermal-vacuum testing campaigns include detector alignment, encircled energy, read noise, dark current, flat field uniformity, detector gain, detector linearity, detector crosstalk, image stability, grism dispersion, filter ghosts, system and filter throughputs, internal calibration system flux level and uniformity, and IR thermal background measurements.

4. Testing Results

The high-level results of these tests include: 1) the first integrated operation of the IR channel; 2) the same good performance of the UVIS and IR detectors seen in previous unit tests before integration into the instrument; 3) the demonstration of routine science operations in flight-like conditions; 4) good margins on the achievable UVIS and IR detector temperatures; 5) lower than expected IR thermal background, based on previous subsystem tests; and 6) excellent optical performance, with UVIS and IR image quality at or near specifications at all wavelengths.

All data analysis results to date have been documented in a series of over thirty Instrument Science Reports (ISRs), which are available for viewing and downloading from the HST WFC3 web site at http://www.stsci.edu/hst/wfc3/documents/ISRs.

4.1. Detector Characteristics

The results of specific tests of detector characteristics are shown in Table 1 and Table 2 for the UVIS and IR channels, respectively. Included in these tables is the Contract End Item (CEI) specification for each characteristic. As can be seen from the tables, the only detector characteristic that does not currently meet the specifications is the IR read noise (note that the values quoted are for a subtracted pair of non-destructive readouts).

Table 1: UVIS Detector Characteristics

Item	Measured Value	CEI Specification
Dark current	0.3 e^-/pix/hour	<20 e^-/pix/hour
Read noise	3.0 e^-/pix	<4 e^-/pix
Linearity	5% deviation at 67,000 e^-/pix	<5% up to 50,000 e^-/pix
Full well	~70,000 e^-/pix	>50,000 e^-/pix

Table 2: IR Detector Characteristics

Item	Measured Value	CEI Specification
Dark current	0.15 e^-/pix/sec	<0.4 e^-/pix/sec
Read noise	~22 e^-/pix rms	15 e^-/pix rms
Linearity	5% deviation at 93,000 e^-/pix	<5% up to 70,000 e^-/pix
Full well	~105,000 e^-/pix	>100,000 e^-/pix

4.2. Optical Characteristics

The results of encircled energy measurements are shown in Table 3.

Table 3: Encircled Energy

Channel	Wavelength	Encircled Energy
UVIS	250 nm	72% in 0.20 arcsecond radius
	633 nm	79% in 0.25 arcsecond radius
IR	1.0 μm	60% in 0.25 arcsecond radius
		75% in 0.37 arcsecond radius
IR	1.6 μm	46% in 0.25 arcsecond radius
		80% in 0.60 arcsecond radius

The blue-optimized CCD's of the WFC3 UVIS channel result in system throughputs that are below that of ACS/WFC at the red end of the optical range, yet far exceed that of WFPC-2 at near-UV wavelengths. Figure 1 shows the througput of the UVIS channel (optics+detector) without filters in place.

WFC3 IR channel throughputs are 50% or more higher than that of the HST NICMOS cameras over the 0.9—1.7μm wavelength range. Figure 2 shows the throughput of the entire IR channel as seen through each IR filter.

5. Pipeline Processing

The `calwf3` pipeline that will be used in the STScI OPUS system to calibrate WFC3 data will be very similar to the `calacs` pipeline in high-level structure (see Hack 1999a). There will be two main branches: one for UVIS channel images and another for IR images. The steps applied on the UVIS branch will be the same as what's used for ACS/WFC images and the IR steps will be similar to NICMOS processing (Hack 1999b; Bushouse, Skinner & MacKenty 1997). Figure 3 shows the high-level data flow through `calwf3` processing.

Basic calibration will include the usual necessary steps such as bias and dark subtraction, flat fielding, and saturation and bad pixel flagging. UVIS images will also have shutter shading and post-flash corrections, as necessary. IR exposures will receive a non-linearity correction and "up the ramp" fitting, which includes CR rejection. Scientific operations of the instrument will allow for subarray readouts in both the UVIS and IR channels, as well as on-chip binning, by factors of 2 and 3, in the UVIS channel. The calibration pipeline will accomodate all of these modes, using subarray extractions of reference images and binned reference images. A major new feature of the WFC3 IR channel detector is the inclusion of a set of "reference" pixels around the perimeter of the the detector array, which are not sensitive to light but are otherwise included in the remaining electronic readout chain.

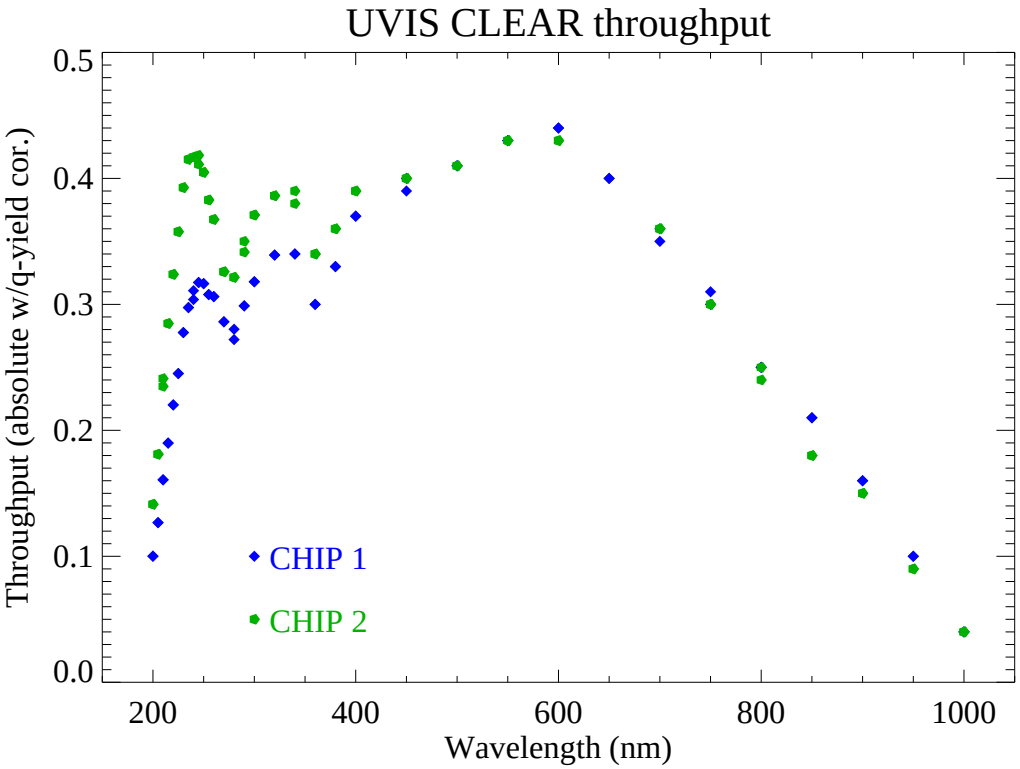

Figure 1: UVIS channel clear throughput

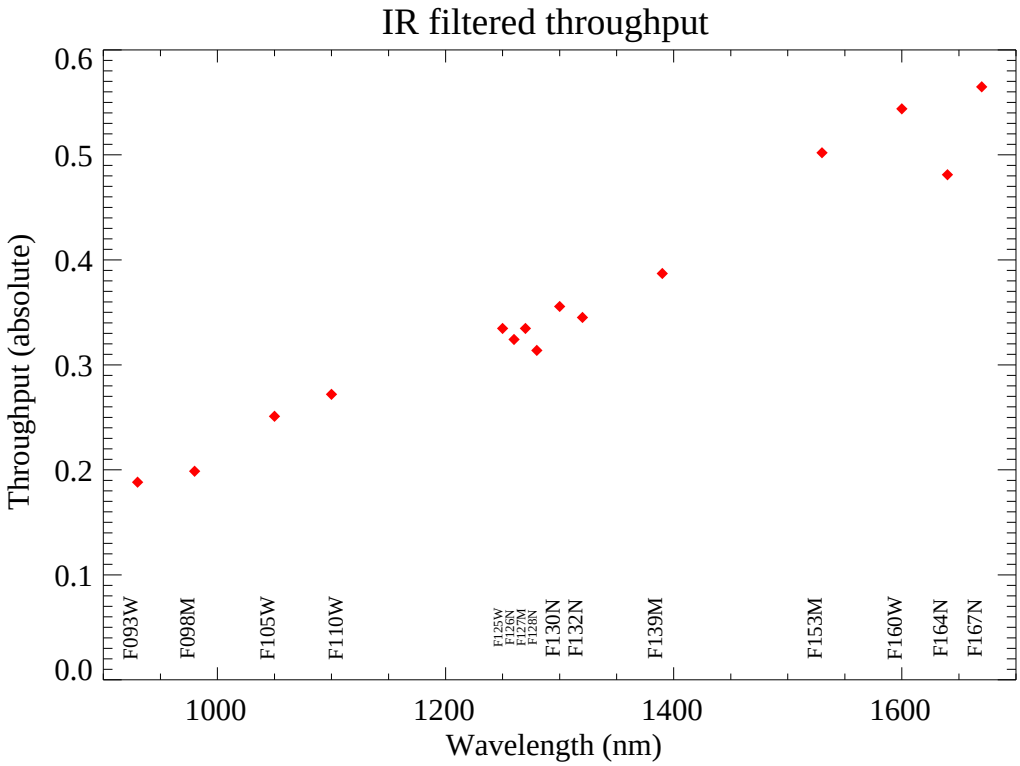

Figure 2: IR channel filtered throughput

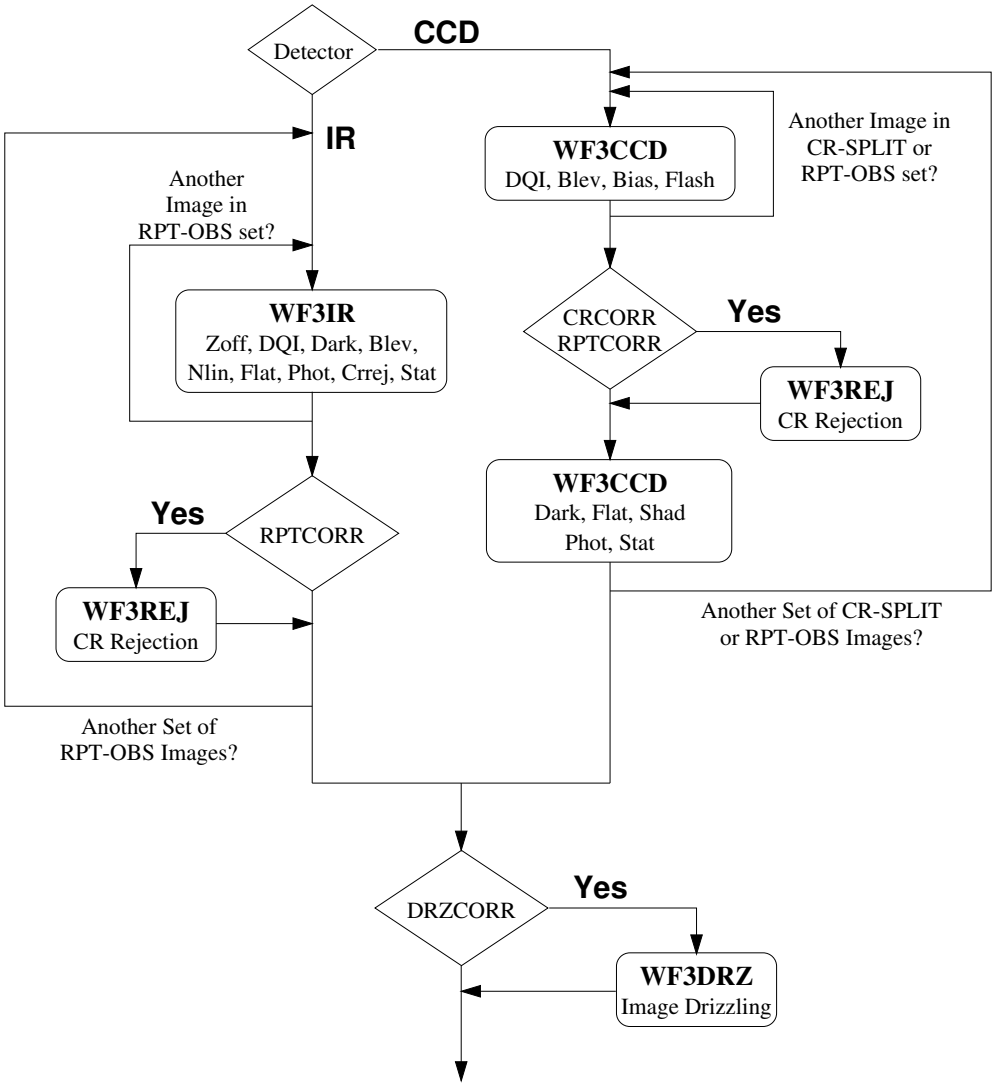

Figure 3: CALWF3 high-level flow

These pixels will be used by the IR branch of `calwf3` to track and remove drifts in the detector bias level from readout to readout within a given exposure.

The ground system will allow observers to obtain multiple exposures on a target, in either CR-SPLIT or REPEAT-OBS modes. These multiple images will be associated and combined during `calwf3` processing. Both UVIS and IR calibrated images will receive drizzle processing, either as individual images or as associated sets (e.g. for dithered patterns). Drizzle processing is necessary, even for individual images, in order to remove geometric distortions and to correct for distortion-induced photometric errors.

References

Bushouse, H., Skinner, C., & MacKenty, J. 1997, *Instrument Science Report* NICMOS 97-28, (Baltimore: STScI)

Hack, W. 1999a, *Instrument Science Report* ACS 99-03, (Baltimore: STScI)

Hack, W. 1999b, *Instrument Science Report* ACS 99-08, (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Characterization Tests of WFC3 Filters

S. Baggett[1], R. Boucarut[2], R. Telfer[3], J. Kim Quijano[1], M. Quijada[2], P. Arsenovic[2], T. Brown[1], M. Dailey[2], D. Figer[1], G. Hartig, B. Hilbert[1], R. Kimble[2], O. Lupie[1], J. MacKenty[1], T. Madison, M. Robberto[1], S. Rice[2], J. Shu[2], J. Townsend[2]

Abstract.

The WFC3 instrument to be installed on HST during the next servicing mission consists of a UVIS and an IR channel. Each channel is allocated its own complement of filters: 48 elements for the UVIS (42 filters, 5 quads, and 1 UV grism) and 17 slots for the IR (15 filters and 2 grisms). While a majority of the UVIS filters exhibit excellent performance consistent with or exceeding expectations, a subset show significant filter ghosts. Procurement of improved replacement filters is in progress and a summary of the characterization tests being performed on the new filters is presented. In the IR channel, while no filter ghosting was detected in any of the filters during thermal vacuum testing, the grisms were found to be installed incorrectly; they have been removed and will be reinstalled. In addition, due to the significantly improved response blueward of 800nm expected in the new substrate-removed IR detector (see Invited talk by R.A.Kimble, this volume), two IR filters originally constructed on a fused silica substrate are being remade using an IR transmitting color glass to block any visible light transmission. Tests of the new IR filters and preparations for the grism reinstallation are summarized.

1. Introduction

The replacement filters being manufactured for WFC3 are subjected to a rigorous set of characterization tests. The test suite is designed to verify conformance of the new filters to the filter requirements. Testing begins at the manufacturer, where numerous checks are performed before shipment of the filters to GSFC. Upon receipt at GSFC, the filters undergo further extensive testing, including a detailed visual inspection, imaging tests (point source and flat field), spectral scans, and environmental checks. Results are analyzed and discussed within the WFC3 team; recommendations for potential replacement filters are presented to the WFC3 Science Oversight Committee for their evaluation and approval. The WFC3 replacement filter characterization process is summarized in Figure 1; the individual steps are described in more detail in the sections that follow.

2. Characterization Tests

Vendor Tests Comprehensive checks are performed by the manufacturer prior to shipment of the filter to the customer. These tests verify the filters' compliance with the specifications and the robustness of the filter coating. Key items measured on the candidate flight pieces

[1]Space Telescope Science Institute

[2]NASA/Goddard Space Flight Center

[3]Orbital Sciences Corp., NASA/Goddard Space Flight Center

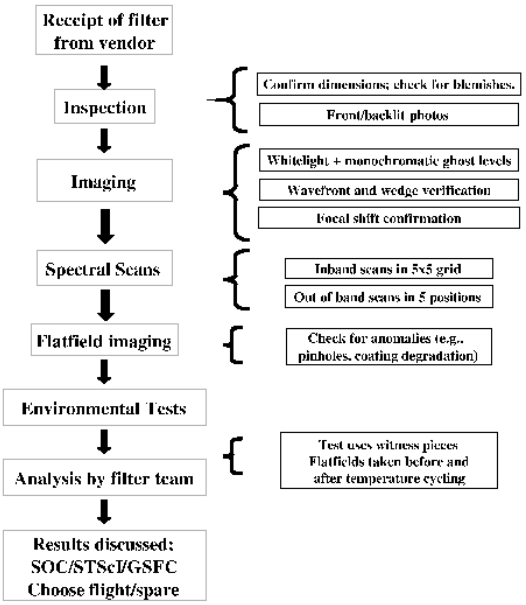

Figure 1: The characterization test process.

include the spectral throughput (primary bandpass as well as out of band blocking), transmissive wavefront (on uncoated substrate as well as the coated filter), and quality control (check of length, width, thickness, inspection for scratches and cosmetic defects, painting of any pinholes, and cleaning). Characteristics such as the coating integrity (adhesion quality, resistance to abrasions, temperature and humidity cycling) and radiation hardness are performed on witness or flight-like pieces.

Inspection Upon receipt of the filter from the vendor, a detailed inspection is performed. The dimensions of the filter are confirmed and the item is inspected for any blemishes or coating defects. Photographs, both front and backlit, are taken to document the condition of the filter and map any pinholes which may be present (see Figure 2a).

Imaging Extensive point source imaging is done with each filter; tests are performed using white (either xenon or deuterium lamp) as well as monochromatic (200-1000nm) light. These data allow for measurement of filter ghosts (strength, appearance, and behavior), as well as wavefront, wedge, and focus shift. Figure 2b shows a whitelight image of a candidate F689M filter. The image was taken at a relatively high filter tilt angle; this moves the ghost away from the PSF and facilitates measurement of the ghost level. The setup ghost is an artifact of the lab setup and appears even without a filter in the beam. Images are taken at a variety of tilt angles and translation positions as well as in configurations simulating seventeen points distributed across the WFC3 field of view. Figure 3a presents a monochromatic scan of the ghosts measured in a candidate filter. The monochromatic results are used to constrain the model, which in turn is used to predict ghost behavior as a function of target spectral type. Figure 3b shows a small subset of images from a phase retrieval dataset. Focus sweeps are obtained through a 9x9 array of 10mm apertures covering the filter, separated by steps of 5mm. Wavefront data taken without a filter are subtracted, including focus and tilt, yielding wavefront error, focus shift, and wedge.

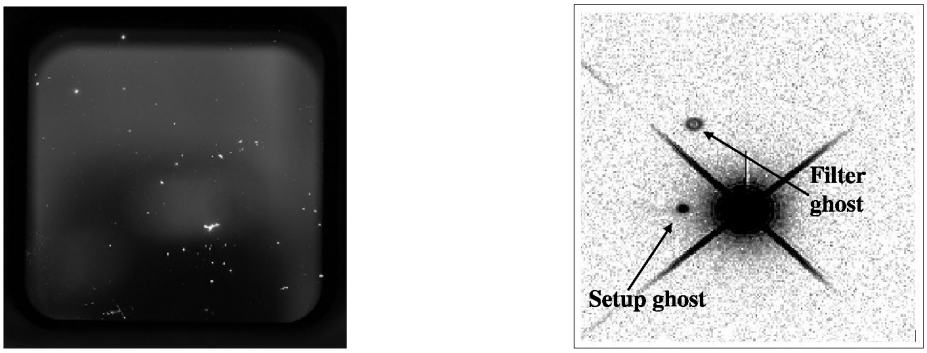

Figure 2: (a) At left is a pinhole image of F280N. (b) At right is an image of the ghost in F689M; filter ghost is at upper left; the feature to the left of the PSF is due to the lab setup.)

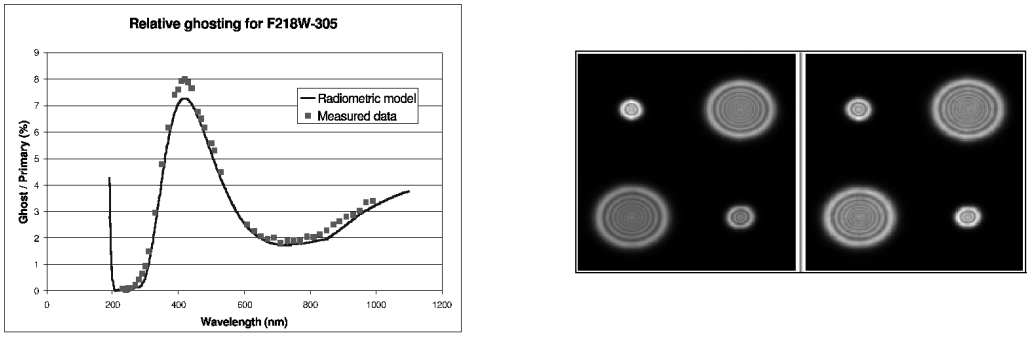

Figure 3: (a) The plot at left shows the excellent agreement between the monochromatic data and model for an F218W filter. Note: the fairly high ghost intensities shown occur out of band, where the filter blocking is significant; integrated over a continuum source, the relative ghost strengths are much lower. (b) The image at right shows a mosaic of a small subset of PSFs taken during a typical phase retrieval test.

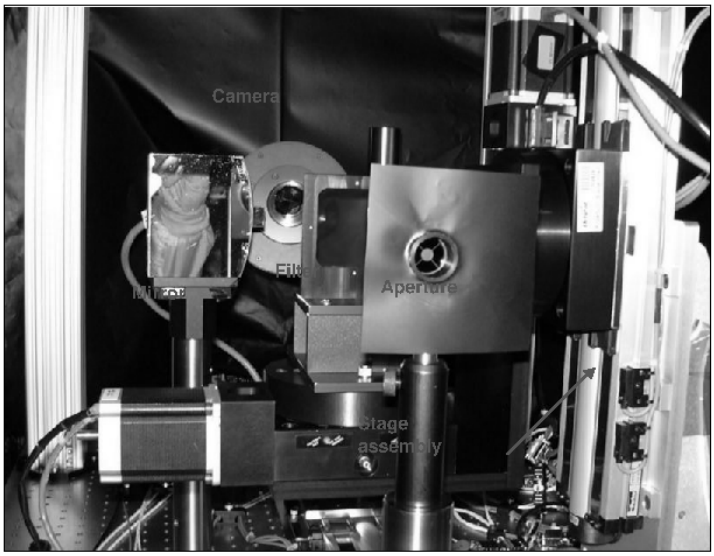

Figure 4: The FilterGEISt lab setup is used for obtaining imaging data of candidate filters.

Imaging tests are performed using FilterGEISt (filter ghost evaluation imaging station, see photo in Figure 4), a dedicated lab setup which reproduces the WFC3 UVIS + HST optical parameters. A 4-axis stage assembly provides repeatable translation/rotation motions as well as simulations of sources within the WFC3 FOV.

Spectral Scans The spectral throughput of the filter is measured both in and out of band using a PerkinElmer 950. Scans are taken at multiple points to allow an assessment of uniformity. In-band UVIS scans are performed in a 5x5 grid across the filter while out of band UVIS scans are executed in the center of the filter and at the four corners. Figures 5a and 5b illustrate typical in and out of band scans which are evaluated in light of the input specifications.

Flat field images The fine scale transmission properties of the filter are verified via flat fielding. The imaging setup is designed to simulate that of the WFC3 internal calibration system (divergent beam). The entire filter is covered, to check for blemishes, coating degradation, or any other anomalies that may be present.

Environmental Tests Filter witnesses and non-flight pieces are subjected to temperature cycling under thermal vacuum (TV) conditions in order to verify the robustness of the coatings. Baseline inspections, spectral scans, and flat fields are taken prior to the test. The cycles are designed to mimic the TV conditions the filters are expected to experience during ground testing and on-orbit. After the cycles, another round of inspections, spectral scans, and flatfields are taken to check for any evidence of coating degradation.

Filter Selection The final choice of which, if any, replacement filters to fly resides with the WFC3 Science Oversight Committee, with analysis inputs from the STScI and GSFC WFC3 teams. The decision is based upon a careful balance of the characterization test results for

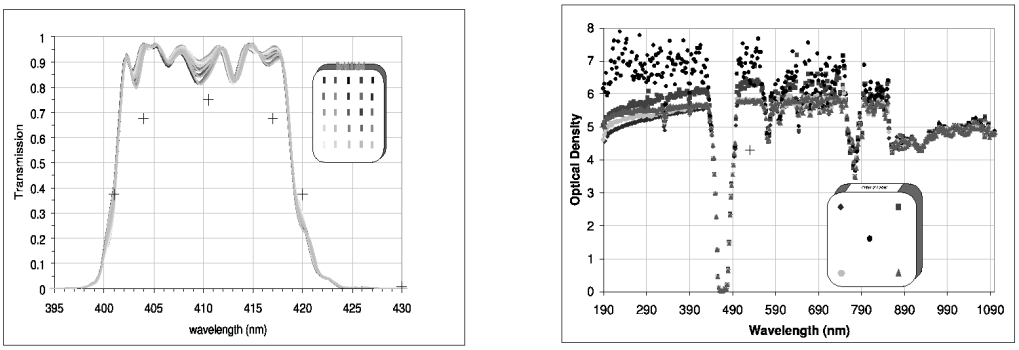

Figure 5: (a) The plot at left presents a representative inband transmission scan, in this case, for an F410M candidate. (b) The image at right shows the blocking scan for an F467M candidate.

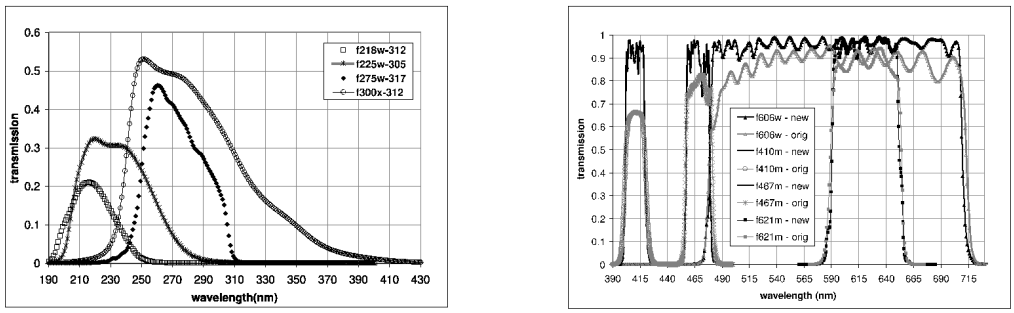

Figure 6: (a) The plot at left summarizes the new UV replacement filter bandpasses. (b) At right are spectral scans for the old and new visible filters.

the new filters versus the characteristics of the current flight filters as determined from previous filter tests as well as instrument-level ambient and thermal vacuum tests.

3. Replacement Filters

The UV is a critical part of the WFC3 bandpass coverage; Figure 6a summarizes the new UV replacement filter bandpasses. In most cases, the filter development technology now available has provided increased throughput in addition to reducing the ghost levels. For example, in F225W (which showed the most improvement) the ghost levels are reduced from about 15 percent to less than 0.3 percent and the new filter grasp (integral of the filter and WFC3 transmission over wavelength in the primary bandpass) is about 30 percent more than the original flight filter. Ghost levels in the visible filters have been reduced as well (some even eliminated) and again, in most cases the grasp of the new filter is higher than in the old filter. Figure 6b provides a comparison of the bandpasses of the original and new visible filters.

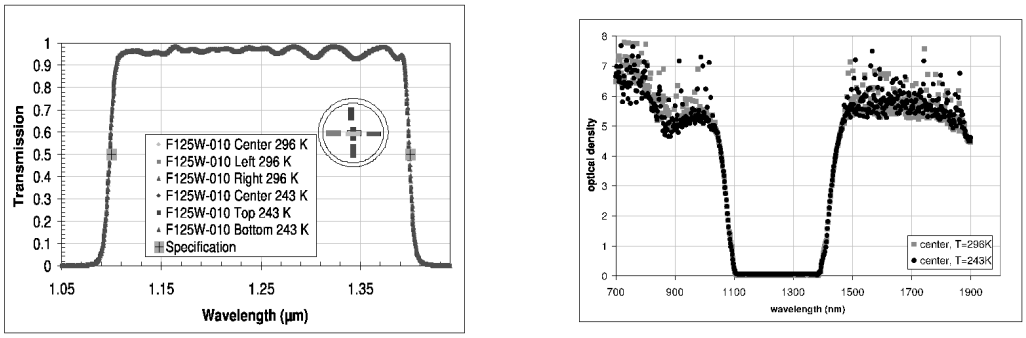

Figure 7: Plots showing typical inband (a) and out of band (b) transmission curves for a candidate F125W filter.

4. IR filter testing

No filter ghosts were found in the IR filters during instrument-level testing but two filters originally constructed on fused silica (because the QE of the original IR detector was low in the blue) must be remade on an IR-transmitting color glass (IR80). The change is driven by the new IR detectors under development, which are being thinned to reduce particle-induced luminescence in the CdZnTe substrate; the new detectors will have significant QE well blueward of 800nm. As with the UVIS filters, the IR replacement filters undergo a full set of characterization tests, including spectral scans at a variety of locations on the filter as well as at two temperatures (room temperature and nominal operating temperature of -30C). Given the relatively small filter size (25.4mm in diameter for IR, 53.2mm square for UVIS), the in-band IR scans are performed at 5 points and the out of band scans at filter center, as illustrated in Figure 7a and 7b. An IR InGaAS IR array camera will be incorporated into the FilterGEISt lab setup, in order to measure wavefront, wedge, and perform a ghost check. Finally, the two IR grisms which were installed incorrectly have been removed and will be reinstalled in the upcoming year. One of the grisms is being remade with an improved red edge in order to reduce the transmission of thermal background radiation, which led to an undesirably high background on the detector in thermal vacuum testing.

5. For more information about WFC3

http://www.stsci.edu/hst/wfc3/ and http://wfc3.gsfc.nasa.gov

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

COS: NUV and FUV Detector Flat Field Status

Steven V. Penton
Center for Astrophysics and Space Astronomy, University of Colorado, Boulder, CO 80309

Abstract. There are two main requirements for the flat fields (FF) of the two (NUV and FUV) COS detectors. The first requirement is that the pixel-to-pixel variations (the P-Flats) of the detectors are uniform to ±20% for each of the two FUV detector segments, and ±6% for the NUV detector. The second requirement is that the derived FFs are able to support signal-to-noise (S/N) observations of 30 under routine observations, and are of sufficient quality to support S/N = 100 data for special observations. As illustrated here, the COS FFs obtained during ground calibration (October 2003) meet these requirements.

1. Introduction

The Hubble Space Telescope's (HST) Cosmic Origins Spectrograph (COS) Far Ultraviolet (FUV) detector consists of micro-channel plates (MCPs) and time-delay anodes to provide photon counting and imaging capabilities. The COS FUV detector does not contain physical pixels, instead it uses time-delay circuits to compute the photon arrival location (in digital elements (DE)). The FUV resolution element (RE) is 6×12-20 DE, depending on cross-dispersion imaging properties of the grating (G130M, G160M, or G140L). The COS Near Ultraviolet (NUV) detector is the STIS spare MAMA, and has physical 25×25 μm pixels with a RE of 3×3 pixels independent of grating (G185M, G225M, G285, or G240L). During initial COS calibration (October 2003), data was obtained to characterize the flat field of both COS detectors. These results are presented here with additional discussion on the topics of dark counts, achievable signal-to-noise, and additional calibrations (and challenges) required for the FUV channel.

2. Status of the COS NUV Detector Flat Fields

There are no obvious structures in the NUV dark images and the count rates are extremely low. A total of 18221 dark counts (cts) were measured in 6800 seconds (s), or 2.68 cts/s. One pixel (p) contained 3 counts, 152 showed 2 counts (0.01%), and 17914 had a single count (1.7%). The average dark count rate is 2.6×10^{-6} cts/s/p (1 ct/p every 4.5 days) or 0.41 cts/s/cm^2. However, the dark count rate is expected to be higher on orbit.

The NUV flat field (FF) illumination was provided by a combination of observations of the COS internal FF lamp and an external source. The total counts in the raw NUV FF image is 3.7 billion counts during an exposure time of 89250 seconds (24.8 hours). The average counts per pixel (p) in the spectral region is ∼8900 counts (S/N ∼95 per p), or ∼80,000 counts per 3×3 resolution element (RE) (S/N ∼285 per RE). The minimum S/N over any RE was ∼220, while the maximum was ∼340.

The smooth dispersion direction (DD) flat field (FF) profile allowed the illumination to be modeled for each column using low-order polynomials along the DD to construct the final 2-D model of the illumination of the spectral region (the L-Flat, left panel of Figure 1). Dividing the raw FF by the L-Flat gives the pixel-to-pixel variations inherent in the detector

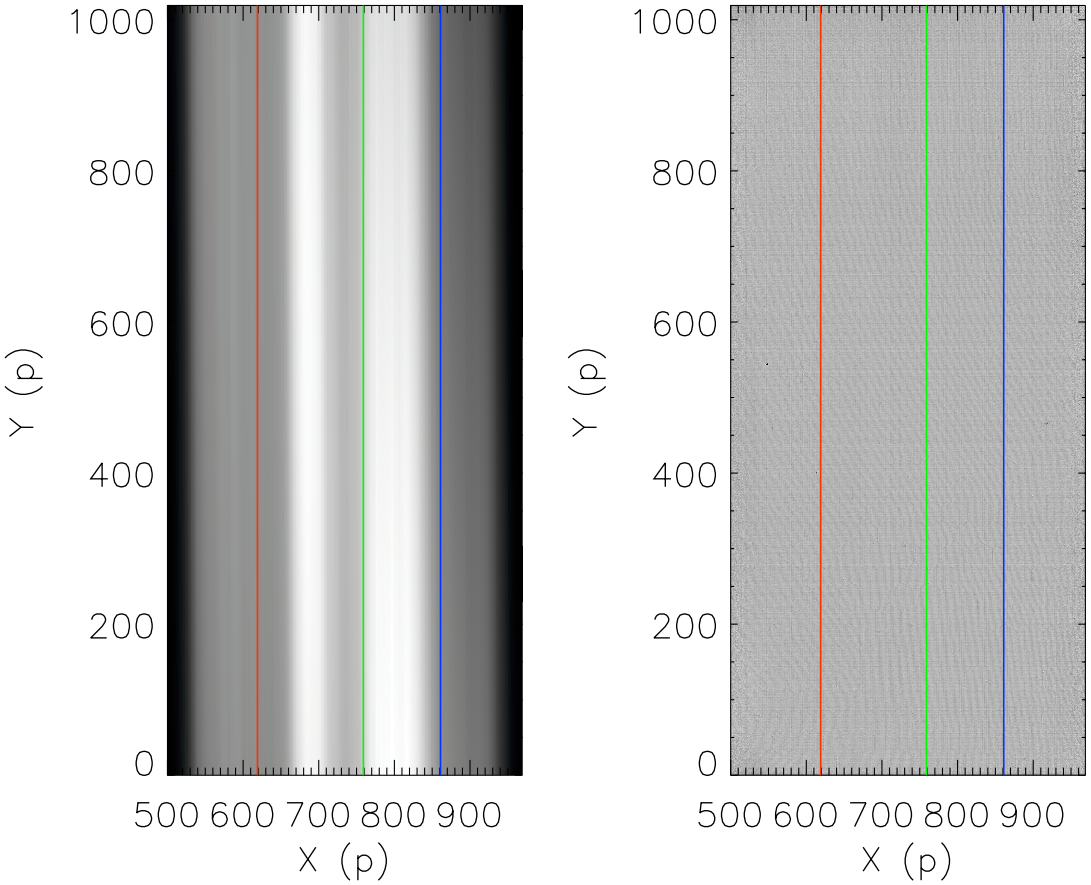

Figure 1: The left panel shows the modelled NUV flat field (FF) illumination (the L-Flat) obtained during COS October 2003 calibration. The right panel shows the COS NUV pixel-to-pixel response function (the P-flat). On both figures, the location of the science spectral stripes at x=620, 760, and 860 are indicated.

(the P-Flat, right panel of Figure 1). The pixel-to-pixel variations in the NUV P-Flat are ±2%, easily exceeding the ±6% requirement.

2.1. Achievable Signal-to-Noise

The pixel-to-pixel variations of the P-Flat determine the absolute maximum signal-to-noise ratio (S/N) achievable without flat field (FF) correction, Figure 2. For the NUV detector, this is a S/N of >50 per RE. Applying the known FF correction to spectral data obtained during initial calibration indicates that by applying a FP_SPLIT (dithering, 4 observations) + FF strategy S/N ∼100 per RE should be achievable with the NUV channel of COS. In fact, the P-Flat variations are so small/uniform that FP_SPLIT data alone can approach this S/N.

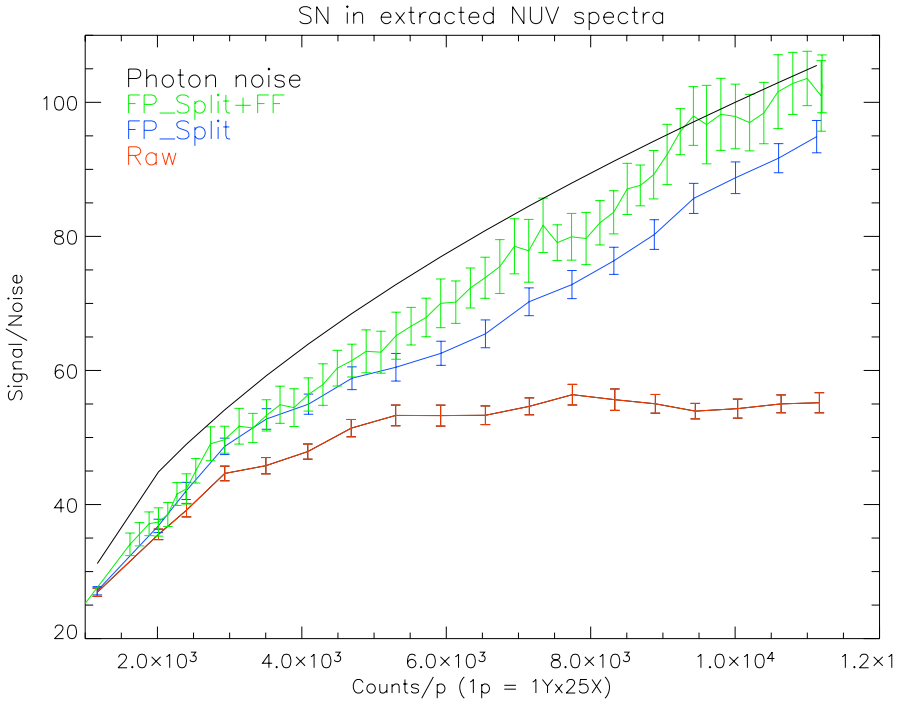

Figure 2: COS NUV achievable signal-to-noise (S/N). A high S/N continuum source, observed during ground calibration, was reduced by a variety of calibration steps (described in the text). A continuum S/N ratio of 30-100 is easily obtained. (S/N decreases in the order of the labels at the upper left.)

3. Status of the COS FUV Detector Flat Fields

All FUV data must be thermally (TC) and geometrically corrected (GC). The accuracy of these corrections are vital to the calibration of COS FUV data. The GC attempts to compensate for non-linearities in the physical location to digital element (DE) mapping, whether caused by the electronics or the MCPs. The thermal correction prevents temporal drifts in this mapping. All FUV data presented hereafter have been TGC corrected.

There are obvious structures in the FUV dark images, but the count rates are extremely low, much lower than for the NUV detector. A total of 532538 dark counts (cts) were measured in 56301 seconds (s), or 9.5 cts/s, or $< 2 \times 10^{-6}$ cts/s per digital element (DE) for segment-A. The count rate for segment-B is similar ($\sim$10 cts/s). As shown in Figure 3, the majority of the dark counts are in hot spots or in non-science areas (upper and lower edges) of the detector. Over the science region of the detector ($450 < Y < 550$), the segment-A count rate is $< 10^{-6}$ cts/s/DE or $\sim$1 ct per DE per fortnight. Hot spots and high-count areas will be filtered during pipeline processing. The dark count images and the flat fields will be used to identify problem areas to be filtered during HST+COS calibration.

The FF illumination, was provided by multiple observations of the COS internal FF lamp. The total counts in each FF image is $\sim$200 million counts. The combined exposure time is 18900 seconds, over 95 sub-exposures (5.2 hours). The average count per digital element (DE) in the spectral region is $\sim$300 counts (S/N $\sim$17.5 per DE), or $\sim$18,000 cts per 5×12 resolution element (RE, S/N $\sim$135 per RE). Figure 4 shows details of the illumination of the segment-A FF.

Piecewise-linear dispersion direction (DD, X) modeling was used to derive the illumination (L-Flat) of the FUV FF. Dividing the FF by the L-Flat produces the pixel-to-pixel variations inherent in the detector (the P-Flat, a portion of which is shown in Figure 5).

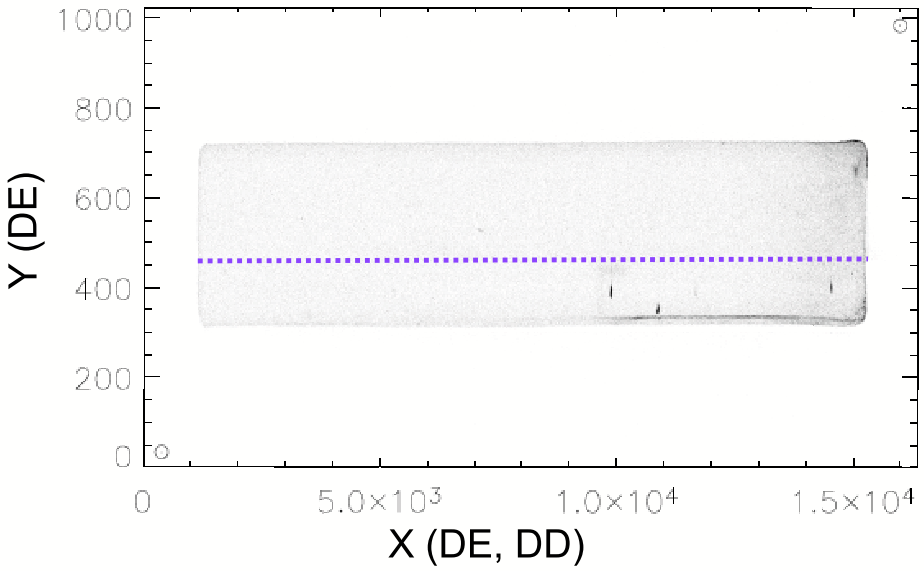

Figure 3: The segment-A dark image shows a count rate of 2×10^{-6} cts/s per digital element (DE). The vast majority of the dark counts are in hot spots or in non-science areas (upper and lower edges) of the detector. In the science region of the detector (dashed line), the count rate is ∼1 ct per DE per fortnight (8×10^{-7} cts/s per DE). Circles in the extreme upper right and lower left indicate the non-physical, digital space, location of the electronic stims in the thermal correction (TC). The dispersion direction (DD) is the X axis.

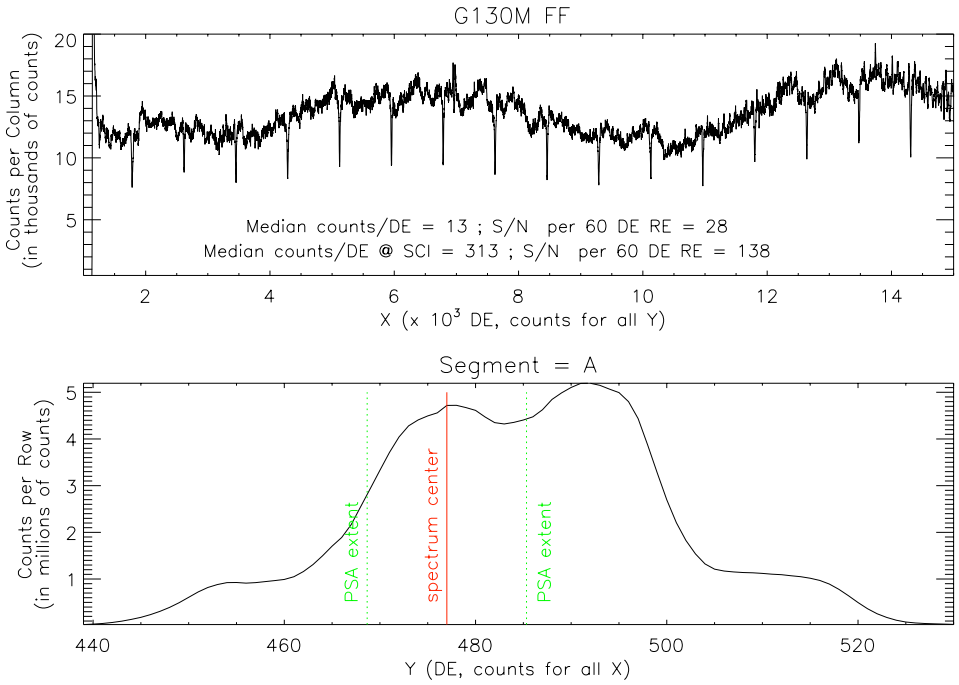

Figure 4: The top figure shows the summed dispersion direction (X, DD) profile of the segment-A FUV raw FF illumination. The DD dropouts are shadows from the detector ion repulsion grid wires. The bottom figure shows the cross-dispersion direction (Y, XD) profile. Also indicated is the location of the spectrum from a point source centered in the primary science aperture (PSA).

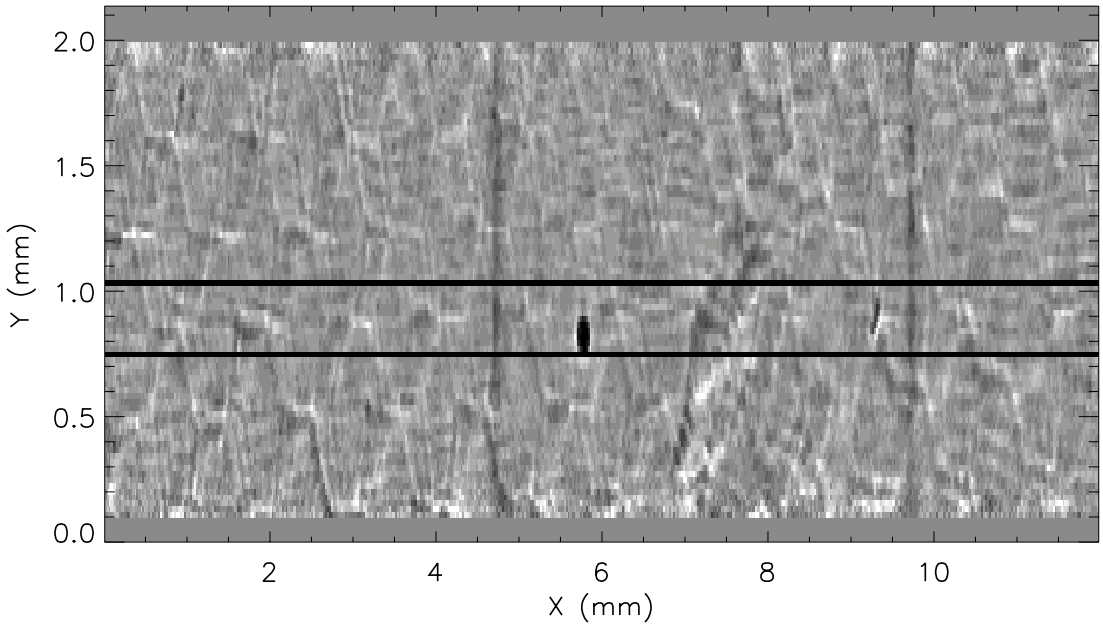

Figure 5: A portion of the FUV segment-A P-Flat is presented in geometrically corrected (GC) physical units. The spectral region is indicated by the horizontal lines. Dead spots, hot spots, plate defects, grid wire, and Moiré patterns are all present, along with the inherent hexagonal micro-channel plate structure (the 'hex' or chicken-wire pattern).

The region displayed in Figure 5 was selected to display the grid wires, and the largest plate defect on the detector. The FUV L-Flat and P-Flats are of lower quality than the NUV flats due to lower count rates per DE, and the inherent structures present in the FUV detector segments The pixel-to-pixel variations in the FUV P-Flat are $\pm 19.4\%$ for segment-A, and $\pm 15.3\%$ for segment-B, meeting the $\pm 20\%$ requirement. Improved techniques and FF photon statistics (obtained during final ground calibration) will only improve our knowledge of the detector segment P-Flats. Residual grid wire and Moiré pattern structures are present in the L-Flat, indicating that further improvement in characterizing the FUV FF is possible.

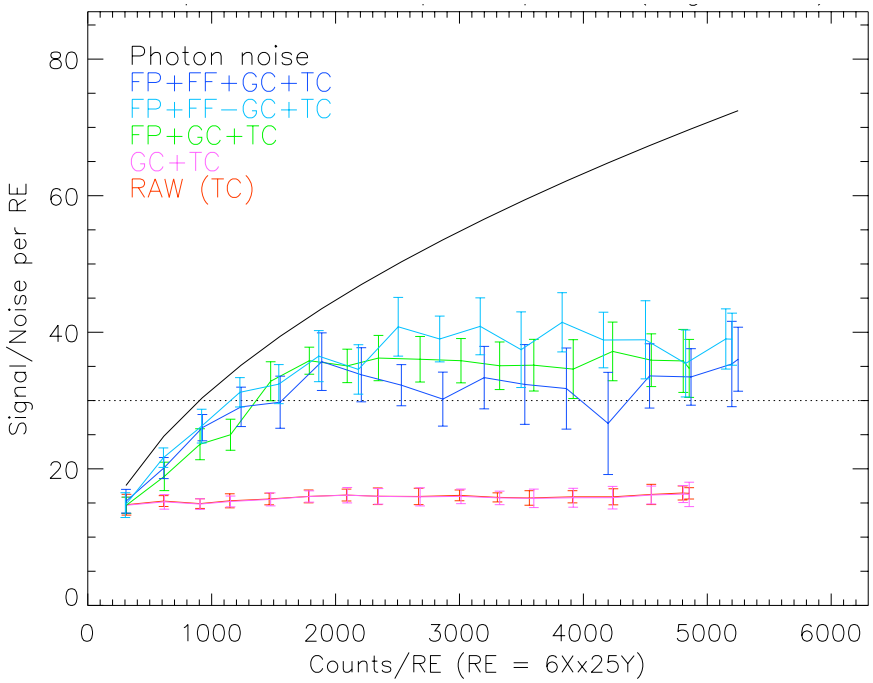

Figure 6: COS FUV achievable signal-to-noise. A high S/N continuum source, observed during ground calibration, was reduced by a variety of calibration steps (described in the text). RAW (TC) and GC+TC both give an S/N ratio ∼15. A continuum S/N ratio of 30-40 is easily obtained with the FP_SPLIT strategy.

3.1. Achievable Signal-to-Noise

The pixel-to-pixel variations of the P-Flat determine the maximum signal-to-noise ratio (S/N) achievable without flat field (FF) correction, Figure 6. This is currently a FUV S/N of ∼5.2 (A) or ∼6.5 (B) per digital element (DE), depending on the segment (A or B). When calculated per resolution element (RE), S/N of 17-20 per RE should be achievable without FF correction. When combined with an FP_SPLIT (FP, ×4) strategy, S/N per RE of 34-40 should be achievable. S/N measurements, taken from the continuum of a high S/N spectrum, after various combinations of calibration procedures are presented in Figure 6 versus counts/RE. Without an FP or FF, the S/N is limited to ∼17. With the FP the S/N increases as expected until S/N of ∼34. Inclusion of a FF correction does not increase the S/N. This is true whether the FF is applied before the GC (FP+FF-GC+TC) or after (FP+FF+GC+TC), however there is some indication (at the 1 σ level) that applying the FF before the GC correction improves the measured S/N. It is unclear at this time whether this is an expected result, or a result of a misalignment between the FF data and the science data, which were obtained through different apertures. Without physical pixels, the FF relies upon the quality of the geometric correction (GC) to turn what would be a normal FF situation (division to correct for a variable response function) into a more complicated deconvolution of the response function.

The FUV channel has five overlapping central wavelength settings, each with an FP_SPLIT (×4) mode. Observing with all settings will provide 20 FP_SPLIT observations which will boost the achievable S/N of ∼40 demonstrated here to ∼100.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

HST Observations in Support of JWST Calibration

James E. Rhoads[1]

Space Telescope Science Institute

Abstract. The James Webb Space Telescope will be NASA's flagship observatory in the next decade. Optimum science return on JWST will require high quality calibration. The Hubble Space Telescope offers unique capabilities that will be essential in achieving such a calibration of JWST. We outline a three-pronged approach for establishing JWST calibration targets using HST. First, ACS astrometry of a dense stellar field in the Large Magellanic Cloud will yield a set of faint astrometric standards that will always be in the JWST field of regard, allowing astrometric calibration whenever it may be needed. Such calibration will certainly be required during on-orbit shakedown, and may also be needed after JWST's periodic primary mirror figure adjustments. Second, new NICMOS spectrophotometry of K giants and main sequence A stars will be combined with existing NICMOS observations of DA white dwarfs and solar analogs to establish a diverse set of sepctrophotometric calibrator stars. These stars will serve as the basic flux calibrators for the JWST NIRSPEC instrument. Finally, NICMOS photometry of solar analog stars in nearby clusters will allow rapid photometric calibration of the JWST NIRCAM shortly after launch. The overall program is currently in implementation, using 19 orbits of HST time.

1. The James Webb Space Telescope

The James Webb Space Telescope (hereafter JWST) is an infrared-optimized, 6.5m diameter telescope with a planned launch date sometime in the next decade. JWST has an effective collecting area of 25m^2. The primary mirror is articulated to fit inside the launch vehicle. This requires a segmented mirror design which unfolds during early on-orbit checkout. The mirror figure can be adjusted by moving the mirror segments independently, and such adjustments may be carried out monthly or even more frequently depending on the realized stability of the mirror figure in orbit.

JWST will have four science instruments:

- NIRCAM, the Near Infrared Camera. NIRCAM is an imaging camera with wavelength coverage $0.7\mu m \leq \lambda \leq 5\mu m$. NIRCAM has two redundant modules, each covering a $2.2' \times 2.2'$ field of view. The fields of the two modules are approximately adjacent on the sky. Each module has two channels, a short wavelength channel ($\lambda < 2.5\mu m$) and a long wavelength channel ($\lambda > 2.5\mu m$). The channels are split by a dichroic, so they share the same field of view. NIRCAM has a selection of broad, intermediate, and narrow band filters, with resolving powers $5 \leq R \leq 100$. The NIRCAM short wavelength channel is critically sampled at $\approx 2\mu m$, and the long wavelength channel at $\approx 4\mu m$. NIRCAM has coronagraphic capability.

[1]Present address: Arizona State University

- NIRSPEC, the Near Infrared Spectrograph. NIRSPEC provides spectra over the range $0.7\mu m \leq \lambda \leq 5\mu m$. There are three separate dispersing element options, with resolving powers $R = 100$, $R = 1000$, and $R = 3000$. NIRSPEC has a microshutter array allowing simultaneous observations of many targets. Additionally, there are fixed slit and IFU options. The NIRSPEC field of view is about 3'.

- MIRI, the Mid Infrared Instrument. MIRI is a very flexible instrument, offering $5\mu m < \lambda < 27\mu m$ imaging, $5\mu m < \lambda < 14\mu m$ spectroscopy with R 100, and $5\mu m < \lambda < 27\mu m$ integral field unit spectroscopy with R 3000. The imaging channel offers coronagraphic capability in four channels, each optimized for a separate wavelength.

- FGS-TF, the Tunable Filter module on the Canadian Fine Guidance Sensor. The tunable filter module provides tunable narrowband imaging with $R \sim 100$ over a $2.3' \times 2.3'$ field, with a wavelength range $\sim 2\mu m < \lambda < 4.8\mu m$. It is critically sampled at $\approx 4\mu m$.

JWST poses some unique calibration challenges. For (spectro)photometry: JWST's wavelength coverage is very broad, and little of it is accessible from the ground. Saturation comes early (especially for sources that can be well observed in the mid-infrared with Spitzer).

For Astrometry: JWST has a "floppy" architecture: its astrometric solution may change comparatively quickly. Moreover, JWST has tight roll angle constraints. Combined, these factors mean that self-calibration of an astrometric field may not be possible with JWST: The best astrometric calibrations depend on observations taken at roughly perpendicular roll angles. Obtaining data at two such roll angles would require $\gtrsim 60$ days delay between epochs with JWST, and in that time, the geometry of the JWST detectors may have changed appreciably.

2. Astrometric Calibration Plans

JWST Astrometric Calibration Requirements:

- Field Distortion must be corrected to 5 mas accuracy, for all Science Instruments and the guider.

- Pixel scale for NIRCAM = 0.032 arcsec, i.e., finer than the ACS-WFC.

- There are also requirements on knowing relative placements of instruments in the focal plane. If HST can help address these, it will be with archival data that spans a wide field of view (e.g. GOODS, GEMS, or COSMOS).

JWST Field of Regard: JWST is a passively cooled telescope. A large sun shade prevents sunlight from falling on the telescope and instruments. The orientation of the sun shade with respect to the telescope is fixed. Thus, the requirement that the telescope always remain in shadow translates to restrictions on the JWST Field of Regard (i.e., the portion of sky accessible to the telescope at any given time). The net result is a minimum sun angle of 85°, and a maximum sun angle of 135° (i.e., a 45° exclusion zone around the anti-sun). Since JWST will be in an L2 orbit, the ecliptic poles become the continuous viewing zone; that is, any point within 5 degrees of either ecliptic pole will always be in the field of regard.

Suitability of Existing ACS Fields:

- Arbitrary ACS-WFC fields achieve accuracies of 5 mas (milliarcseconds) easily.

- < 2 mas is achievable if the ACS distortion is itself properly calibrated. The best way to verify this is by observations at multiple roll angles. (See, e.g., Pirzkal et al 2005, for astrometry of stars in the Ultra Deep Field.)
- Existing extragalactic deep fields: The UDF and similar fields are too sparse.
- Existing globular cluster fields are few, none are in the JWST CVZ, and they have other issues- brightness, stellar density gradients.
- Requirements for JWST astrometry fields: Stellar density (high!); Stellar brightness (faint!); Proper motions (small!); and Continuous availability a major plus.
- Conclusion: We need a well calibrated ACS field in the JWST CVZ.

Astrometric Observations Overview: We will establish a new astrometric field in the Large Magellanic Cloud, using 7 orbits of ACS WFC observations. The LMC offers several advantages. The stellar density is high, the field can be in the JWST continuous field of regard, and the distance to the stars is sufficiently large that proper motions between the time of the HST observation and the JWST launch will be negligible (or nearly so).

To get a suitable density of stars at a suitable brightness, it is necessary to go near the peak projected surface density of LMC stars. This is about 4° from the south ecliptic pole. Observations with a JWST integration time of a few minutes will then reach the depth required for optimum astrometry. This depth is just below the crowded field astrometry regime, where stars with measured signal to noise $s/n \gtrsim 100$ have a typical separation of order $\sim 10\times$ larger than the point spread function. The target flux level, $V \sim 23$, is also near the "knee" in the observed LMC luminosity function (Geha et al 1998). This minimizes the confusion effects due to stars too faint to be used as astrometric standards themselves (I.e., the number of stars at $V > 23$ is fewer than the extrapolation of the $V < 23$ number-magnitude relation for the field.)

The field will be observed at two roll angles, separated by about 90°. At the first roll angle, we will use a cross-shaped pattern with an edge-to-edge extent of 2 ACS fields. The second roll angle will cover just the central ACS field of view. Within each orbit, the ACS will be dithered on small to intermediate steps. In net, then, we will have sampled each point in the central ACS field of view with essentially every quadrant of the ACS WFC detectors, and at two widely separated roll angles, and with offsets of intermediate scales too. This will allow a lot of redundancy in solving separately for the celestial coordinates of stars and the distortion terms in the ACS WFC. The ACS WFC distortion information will be used to help achieve our ultimate goal, which is good astrometric data on the star field.

The precise field selected needs to account for the desired stellar density. It also needs to avoid stellar density gradients, dust, and edge of CVZ.

3. Spectrophotometric Calibration Plans

Overview:

- HST uses White Dwarf and Solar Analog calibrators.
- Spitzer/IRAC uses K giants and main sequence A stars.
- Spitzer finds a discrepancy between K and A stars (Reach et al 05).
- Planned observations will allow calibration of JWST using a wide range of spectral types, including primary standards for both HST and Spitzer.

- HST observations of 8 stars from Spitzer's primary standard and candidate standard lists will cross-calibrate HST and Spitzer.
- Including the Spitzer standards will also make the HST/NICMOS calibration more robust by broadening its base.

JWST Photometric Calibration Requirements: There is a requirement that JWST be capable of achieving data calibration into physical units with absolute accuracies shown in table 1.

Instrument	flux % (imaging)	flux % (coronagraphy)	flux % (spectroscopy)	wavelength (% of resolution element, spectroscopy)
NIRCAM	5	5	NA	NA
NIRSPEC	NA	NA	10	12.5
MIRI	5	15	15	10
FGS-TF	5	10	NA	10

Table 1: Photometric calibration requirements for JWST.

Suitability of Existing Spitzer / HST Data:

- Absolute spectrophotometric accuracy of Spitzer IRS is 5 to 10% (at launch + 1 year).
- This is adequate (just) for JWST spectroscopy, but not for photometry.
- Model spectra tied down by a suite of IR to optical (or UV) data can do better.

Overall Strategy for Spectrophotometric Calibration:

- Take calibrated models as "ground truth."
- Use multiple stellar types to overcome systematic uncertainties in models for any one stellar class: White dwarfs, solar analogs, K giants, A dwarfs.
- Use HST, Spitzer, and ground based data to (i) Select the best model for each star; (ii) show that there are no significant discrepancies between the stellar spectrum and this model; (iii) Adjust the instrument model if necessary, then iterate.
- Stellar astrophysics and detector physics suggest separate lists of standards for MIRI and NIRSPEC. All stars have nearly Rayleigh-Jeans spectra in the mid-IR, so that MIRI fluxes will be much lower than NIRSPEC fluxes. Moreover, MIRI will have a higher sky background, since zodiacal dust emits strongly at $\lambda \geq 5\mu m$.
- We want $\sim$ 16 spectrophotometric calibrators in total: (4 stellar classes) * (4 stars per class).
- Why 4 stars per class? To cover variations of T_{eff} , $\log(g)$, abundance patterns, reddening, etc. within each class.
- MIRI calibrators? Used over the wavelength range $5\mu m < \lambda < 28\mu m$. Hence, there is no urgent need for new data at $\lambda < 2.5\mu m$ here.
- FGS-TF calibration? Can use a combination of NIRSpec and NIRCam calibrators; a separate list probably not needed.

Planned Observations for NIRSpec Spectrophotometry:

- We have selected calibration targets from Spitzer IRAC calibrators and HST calibrators. These have suitable fluxes for NIRSpec.
- We need good calibration from optical to 5 microns, not beyond.
- HST NICMOS grism spectra have high priority here.
- Spitzer IRAC (3.5, 4.5, 5.8, 8 micron photometry) is vital too.
- Spitzer MIPS (mid-IR photometry) and IRS (mid-IR spectra) are optional.
- DA white dwarfs and solar analogs: HST spectrophotometric standards furnish 3 of each; no new HST data needed here.
- Calendar coverage about 80-85% with existing WD and solar analogs.
- A dwarf and K giant calibrators: Spitzer IRAC calibrators furnish ¿ 4 of each, in the JWST CVZ.
- NICMOS grism: < One orbit/target gives $s/n > 100$ from $0.8\mu m$ to $2.5\mu m$. So, 8 orbits total for 8 targets.

The selected stars are all from the Spitzer standard stars listed in Reach et al 2005. All are in the JWST continuous viewing zone. They were chosen to have magnitudes $K_s \gtrsim 11$, and (within reason) to span a range of surface temperature for both the A dwarfs and K giants. The A star IDs (Reach et al) are 1740346, 1805292, 1743045, and 1812095. Those of the K giants are KF08T3, KF01T5, KF06T1, and KF06T2.

4. Photometric Calibration Plans

Ultimately, we would like to tie calibration of JWST imagers to the spectrophotometric calibration via JWST spectrophotometry of fainter sources, but that is a long-term process. It begins with calibration of the JWST spectrographs using the spectrophotometric standards described above, continues with using the JWST spectrographs to observe new standards that are too faint for HST and Spitzer, and ends with obtaining photometric observations of those new standards with the JWST imagers. (Note, the JWST imagers saturate on the spectrophotometric standards described above, and will saturate on anything that can be properly observed with the HST and Spitzer IR spectrographs.)

Meantime, we want some calibrated standard star fields that we can use for rapid on-orbit calibration of imagers.

Our program consists of 4 orbits of NICMOS photometry, to observe four Galactic clusters as calibration targets for JWST NIRCam. At least two of these will be observed with Spitzer IRAC, allowing cross-calibration among the three observatories.

Strategy for Photometric Calibration: Advantages of clusters:

- Many targets per field, hence, efficient.
- Clusters have uniform metallicity, age, distance, reddening =¿ Relevant model spectra are a one-parameter family.

Approach:

- Take good optical and IR photometry;

- Determine the spectral class of each cluster member;
- Use model spectra appropriate to cluster members to predict their fluxes in JWST passbands.

Role of HST and Spitzer:

- HST offers high photometric precision and stability.
- Spitzer offers, in addition, mid-IR wavelength coverage.

Existing and Planned Data for Photometric Calibration Clusters: NIRCam candidate calibration clusters include NGC 2420, NGC 2506, NGC 6791, and Melotte 66. These are good solar analog clusters: Their age is comparable to that of the Sun; their distance is large enough that a solar analog will not saturate NIRCAM in a minimal 10 second exposure, and their metallicity is also approximately solar.

Spitzer observations of NGC 6791 and Melotte 66 were planned last spring (George Rieke, private communication). NGC 6791 was observed, while Melotte 66 was not, due to a Spitzer safing event. At least one of these two clusters is always in JWST field of regard.

At present, ACS F606W and F814W data exist for NGC 6791, and WFPC2 F555W and F814W exist for NGC 2420.

Planned observations: Photometry One orbit of NICMOS photometry gets one or more field at $s/n \gg 25$ for Vega magnitude = 22 in F110W and F160W. So 4 orbits total would give a small set of faint standards suitable for early NIRCam calibration in each of these clusters. NIC2 will be the primary camera, due to its combination of adequate sampling and good photometric stability (Roelof De Jong, personal communication). This photometry will also improve NIRCAM-NICMOS cross calibration.

5. Summary

HST observations will play a key role in meeting JWST calibration requirements.

- HST astrometric observations are essential for JWST calibration.
- HST spectrophotometric and photometric observations will establish a stronger basis for JWST calibration.
- An affordable investment of HST time will save JWST observing time and reduce schedule risk in JWST commissioning.
- The planned observations will cross-calibrate HST, Spitzer and JWST; and can help make HST calibration more robust.

The total orbit budget allocated to this work is 19 orbits. Of these, Astrometry is 7 orbits with ACS, Spectrophotometry is 8 orbits with NICMOS, and photometry is 4 orbits with NICMOS. The observation planning is largely complete for astrometry and spectrophotometry, and is in progress for photometry. As the present author has now left STScI, the final stages of photometric calibration planning are being carried forward by Rosa Diaz-Miller and Ralph Bohlin. Observations for all programs should be obtained within the coming year.

Acknowledgments. I would like to thank all of my colleagues whose support has made this project possible. The list of people who helped is long, and includes (among others) Eddie Bergeron, Richard Hook, Keith Noll, Bill Sparks, Bahram Mobasher, Ralph Bohlin, Ron Gilliland, Linda Dressel, Massimo Stiavelli, Mike Regan, Roleof de Jong, Santiago Arribas, Stefano Casertano, Jeff Valenti, Vera Platais, Peter McCullough, Jerry Kriss, Matt McMaster, Jay Anderson, John Stauffer, George Rieke, Bill Reach, and Rosa Diaz-Miller. Particular thanks are due to Jerry Kriss, Stefano Casertano, and Massimo Stiavelli, for their support and advice throughout; to Ralph Bohlin, for his help with spectrophotometric calibration strategy; to Jay Anderson and Vera Platais, for their advice on astrometric calibration strategy; and to Rosa Diaz-Miller and Ralph Bohlin for their ongoing support of this project.

References

Geha, M., et al 1998, AJ 115, 1045.

Pirzkal, N., et al 2005, ApJ 622, 319.

Reach, W. T., et al 2005, PASP 117, 978.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

NIRSpec Pipeline Concept: A High Level Description

Guido De Marchi, Torsten Böker and Peter Jakobsen

Space Science Department, European Space Agency, Noordwijk, Netherlands

Abstract.
NIRSpec is a multi-object spectrograph (up to ~ 100 objects) with a large field of view ($\sim 10'$ square) and a detector response covering nearly a factor of 10 in wavelength ($0.6 - 5\,\mu$m). It will be employed on a segmented off-axis telescope in deep space. The combination of these factors results in a number of unique challenges for the calibration of the NIRSpec instrument. In this paper we outline some of these challenges and present a high-level concept of the NIRSpec pipeline data reduction sequence.

1. Why is NIRSpec so special?

In general, the goal of the spectrophotometric calibration of any astronomical instrument is to convert — for every field point, wavelength, and "observing mode" (i.e. combination of filter, grating and band) — the number of electronic counts registered in the detector electronics (measured in digital numbers or DN) to the signal received from the astronomical source (measured in erg/s/cm^2/Å or similar units). In the case of a multi-object spectrograph (MOS) in general, and NIRSpec on board JWST in particular, there are a number of effects that make this calibration particularly challenging:

1. At variance with a traditional long slit spectrograph, every detector pixel in a MOS must be calibrated for its response over the entire spectral range, because spectra from a large number of celestial objects can illuminate the same detector pixel, albeit at different wavelengths. For a MOS, therefore, the "flat field" response map is a three-dimensional data cube that "stacks" the two-dimensional, ($2\,\mathrm{K} \times 4\,\mathrm{K}$) pixel response maps along the wavelength dimension. In principle, the wavelength axis must be sampled at the spectral resolution of the instrument (up to 2700 in the case of NIRSpec). Possible simplifications of this approach, which would otherwise result in rather large and impractical calibration reference files, must be sought.

2. The two-dimensional NIRSpec field of view (FOV), combined with the use of reflection gratings and the NIRSpec design (see Figure 1), causes light to hit the dispersive elements at rather large (and varying) incidence angles. This leads to a significant curvature of the projected slit apertures on the NIRSpec detectors: while each individual shutter aperture is essentially tilted with respect to the dispersion axis, the tilt angle varies significantly across the FOV. This complicates calibration and data reduction because each shutter spectrum must be rectified differently.

3. Unlike for ground-based MOS, for NIRSpec the telescope PSF is a strong function of wavelength: the FWHM varies nearly by a factor of 10 between the blue and red ends of the NIRSpec spectral range. On the other hand, NIRSpec uses only a single imaging scale for its entire wavelength range, and the Micro-Shutter Array (MSA) offers only a single physical slit width. Although the shutter width (200 mas) has been carefully selected to optimize throughput and resolution across the entire

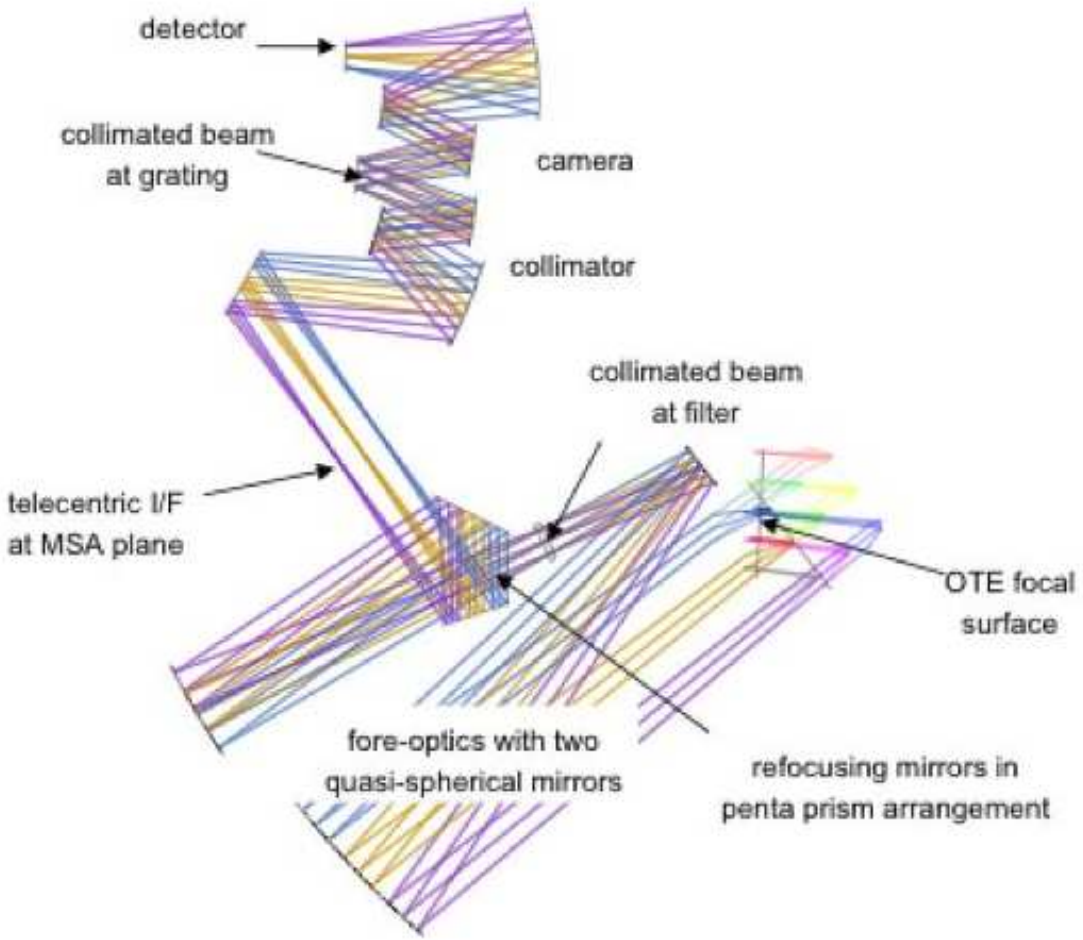

Figure 1: NIRSpec optical layout.

spectral range, the fraction of light that passes through an MSA shutter (or fixed slit) is nevertheless a strong function of wavelength. In principle, the slit throughput can be accurately modeled, but this approach requires exact knowledge of the source shape and position within the shutter. Given the faintness of targets and the limited accuracy of the target acquisition process, such a priori knowledge is not generally available. Note that the slit loss variations with wavelength are most pronounced for a point source, while in the case of a source that uniformly illuminates the entire shutter or slit, a very small correction for chromatic slit loss is needed. Another effect that has consequences similar to those of chromatic slit loss are diffraction losses due to the finite size of the NIRSpec optics behind the MSA, in particular the gratings. The diffraction pattern depends on wavelength, source shape and position within the slit, and thus the amount of light that is diffracted outside the NIRSpec optical train is difficult to predict. Because both these effects — the geometric light loss due to "cutting off" the PSF at the shutter/slit, and the diffraction losses in the optical train behind the MSA — have similar consequences, we refer to them here collectively as "chromatic slit loss".

4. Because of the large field of view, use of a purely reflective optical train, and deployment behind a fast, off-axis telescope, optical distortion is significant in NIRSpec. The resulting plate-scale variations across the aperture provide additional challenges for the spectrophotometric calibration because one must account for the varying solid angle seen by both detector pixels and MSA shutters.

The complexity of the effects discussed above, combined with the limited testing facilities on ground and stringent efficiency requirements in orbit, leads to an approach for NIRSpec calibration which relies heavily on optical modeling of the instrument. Wherever possible, we will attempt to construct parametric descriptions of the instrument features that will be tested, verified, and optimized at every stage from sub-system to instrument level to in flight operations. The requirements for the instrument performance simulator are thus intimately connected to the design of the NIRSpec pipeline.

Here we provide a brief, high level summary of the steps involved in processing NIRSpec spectra. In particular, we address the operations that must be carried out on raw data to render them in a form suitable for scientific analysis.

2. Pipeline goals

The main product of the NIRSpec pipeline is a rectified, finely sampled "window" for each open MSA shutter. This window is a two-dimensional array, sized to capture all the (dispersed) light entering the collimator/camera optics through the shutter under consideration for the specific observing mode. Once processed by the pipeline, all rows in the final extraction window have the same, uniform wavelength scale, i.e. all columns correspond to the same wavelength and the spectral sampling is constant along rows and the same for all extraction windows belonging to a common observing mode. The value of each element (pixel) in the window is the astrophysical flux, calibrated in absolute terms. It should be noted that there is no one-to-one correspondence between array elements in the final extraction window and detector pixels.

One extraction window is produced for each open shutter, regardless whether it contains objects or background, but each window is marked accordingly on the basis of the information provided by the observer. The correct procedure to subtract background light depends on the size and shape of the source, as well as on the structure and spectrum of its surroundings. Therefore, background subtraction from the extracted spectra is left to the observer, who will select the background windows most suitable to his needs. However, for quick-look purposes, the pipeline computes an average background by combining all available background windows, or only those specifically associated by the observer to one target, and subtracts it from the extracted object spectrum. The window is then collapsed in the spatial direction and a one-dimensional spectrum is produced, which is further corrected to compensate for the chromatic slit loss (using the correction computed for a point source). Since, however, the latter correction depends crucially on the geometry of the target and on its exact position within the slit, the produced one-dimensional spectrum should only be used for a first assessment of the overall quality of the data (quick-look). Scientific analysis should be carried out on the full two-dimensional final extraction window.

3. The NIRSpec pipeline

The pipeline is a series of operations that need to be carried out on detector raw data to remove instrumental signatures and to convert counts to absolute flux units. For NIRSpec, and MOS instruments in general, it is convenient to split these operations in two groups: 1) those that depend solely on pixel coordinates on the detector; 2) those that depend on

pixel coordinates, wavelength of the incident light and location of the (undispersed) source in the field of view.

The first group includes dark current, linearity and bias level corrections, all of which are independent of the disperser and filter used, as well as pixel-to-pixel variations in the response of the detector (hereafter "P-flat"), for which the wavelength dependence is expected to be negligible within each passband, as suggested by preliminary analyses. Operations belonging to this group can be applied to the raw data frame as a whole and are not discussed further here because they are all similar to those routinely carried out for the instruments on board the HST.

The second group includes corrections for sensitivity variations (on both global and local scale), geometric distortion (affecting the flux, spatial scale and wavelength of each pixel) and absolute flux calibration. Operations of this type are more logically executed at the level of the extraction window because they depend on which specific MSA shutter is used. These steps are conceptually not different from those normally carried out with traditional ground-based single and multi-object spectrographs, and include: removing large scale variations in the uniformity of the detector's response, both as a function of position and wavelength (usually referred to as L-flat); tracing the spectrum in the aperture and rectifying the wavelength scale and the spatial axis (spectral extraction); assigning a wavelength to each pixel in the extracted spectrum (wavelength calibration); correcting the measured count rate to account for field and wavelength variations in the throughput of the optics and dispersing elements, such as the grating blaze function (flux calibration); registering the measured count rates with those of a reference spectrophotometric standard (absolute calibration).

In reality, however, the complexity of NIRSpec (see Introduction) and the constraints imposed by having the MSA in the optical train have serious repercussions on the pipeline flow and on the calibration of the instrument, both on the ground and in orbit. In particular, some of the steps in the second group indicated above must be carried out simultaneously, since they require calibration measurements that are intertwined and cannot be otherwise obtained separately.

An example of two calibration steps that must be executed together is the correction for large scale flat field uniformity and throughput variations. The former is the so called L-flat, i.e. the measure of the low-frequency variations in the response of the detector, whereas the term 'throughput' here refers to the transmission of all the reflective and dispersive elements along the optical path. Both depend on field angle (or pixel position) and wavelength, and both are characterized during ground calibration of the individual components and at the integration level. This calibration information could in principle be used to determine the total sensitivity as a function of wavelength and field angle for any observing mode. However, while the wavelength dependence can be accurately described by a model, validated by ground calibration measurements, the dependence of the throughput on field angle is much more difficult to predict analytically and will have to be determined empirically. Furthermore, such a determination requires a proper knowledge of the L-flat, which may not be attainable with the required level of accuracy during ground calibration (because, after integration, the detectors cannot be uniformly illuminated due to the presence of the MSA in the optical path).

It is, therefore, advisable to determine simultaneously the combined effect of throughput and detector response on the overall sensitivity variations of the instrument, as a function of wavelength and field position, since they collectively affect the response of each pixel as a function of wavelength. To this purpose, NIRSpec is equipped with an internal continuum calibration lamp, of known spectrum, that uniformly illuminates the MSA. The lamp can be used to simultaneously measure on the detector, during thermal vacuum and in flight, the combined response (i.e. DQE + throughput) as a function of wavelength and position. Time constraints do not permit such measurements to be obtained for each individual shutter, but a suitable subsample may be sufficient. The result would be a series of

extraction windows uniformly spread across the detector, each containing the same known spectrum, which completely cover the detector at least for some wavelengths.

With measurements of this type, it is possible to derive the low-frequency detector response variation, both spatial and spectral, folded with the low frequency spectral and field variation of the throughput. It is useful to visualize this calibration file in the form of a three-dimensional array, with X and Y along the detector's coordinates and Z scaling with wavelength. Since it is expected that any variations in the overall instrument response, with position or wavelength, are very slow and thus suitable to be interpolated via a low-order polynomial, it is only necessary for the pipeline to know the form and parameters of this function in order to apply the correction. The advantage of the approach mentioned above, over the separate measurement of the detector L-flat and individual throughput of the components, is that it can be easily repeated in flight, for a subset of the exposures, to verify the temporal stability of the sensitivity.

4. Conclusions

Although space limitations do not allow us to address in detail the many individual tasks of the NIRSpec pipeline, we offer in Figure 2 a schematic conceptual outline of the steps needed to deal with the MSA data. It is obvious from the complexity of the tasks at stake that, in order to efficiently carry out the calibration of the NIRSpec data, an innovative approach will be needed that departs in more ways than one from that routinely used for HST data. In particular, the pipeline will have to rely more heavily on optical modeling of the instrument and the reference files will, in most cases, be replaced by tables of coefficients that represent the best fitting parameters of those models. Furthermore, the multiplicity of observing modes makes it almost impossible to test and derive reference files for each and every observational setup during the ground campaign. Therefore, developing and validating an end-to-end simulator of the scientific performances of NIRSpec becomes a pressing need in order to pave the road to the smooth and manageable commissioning and calibration of the instrument.

Acknowledgments. We are grateful to Pierre Ferruit for very useful discussions on this topic.

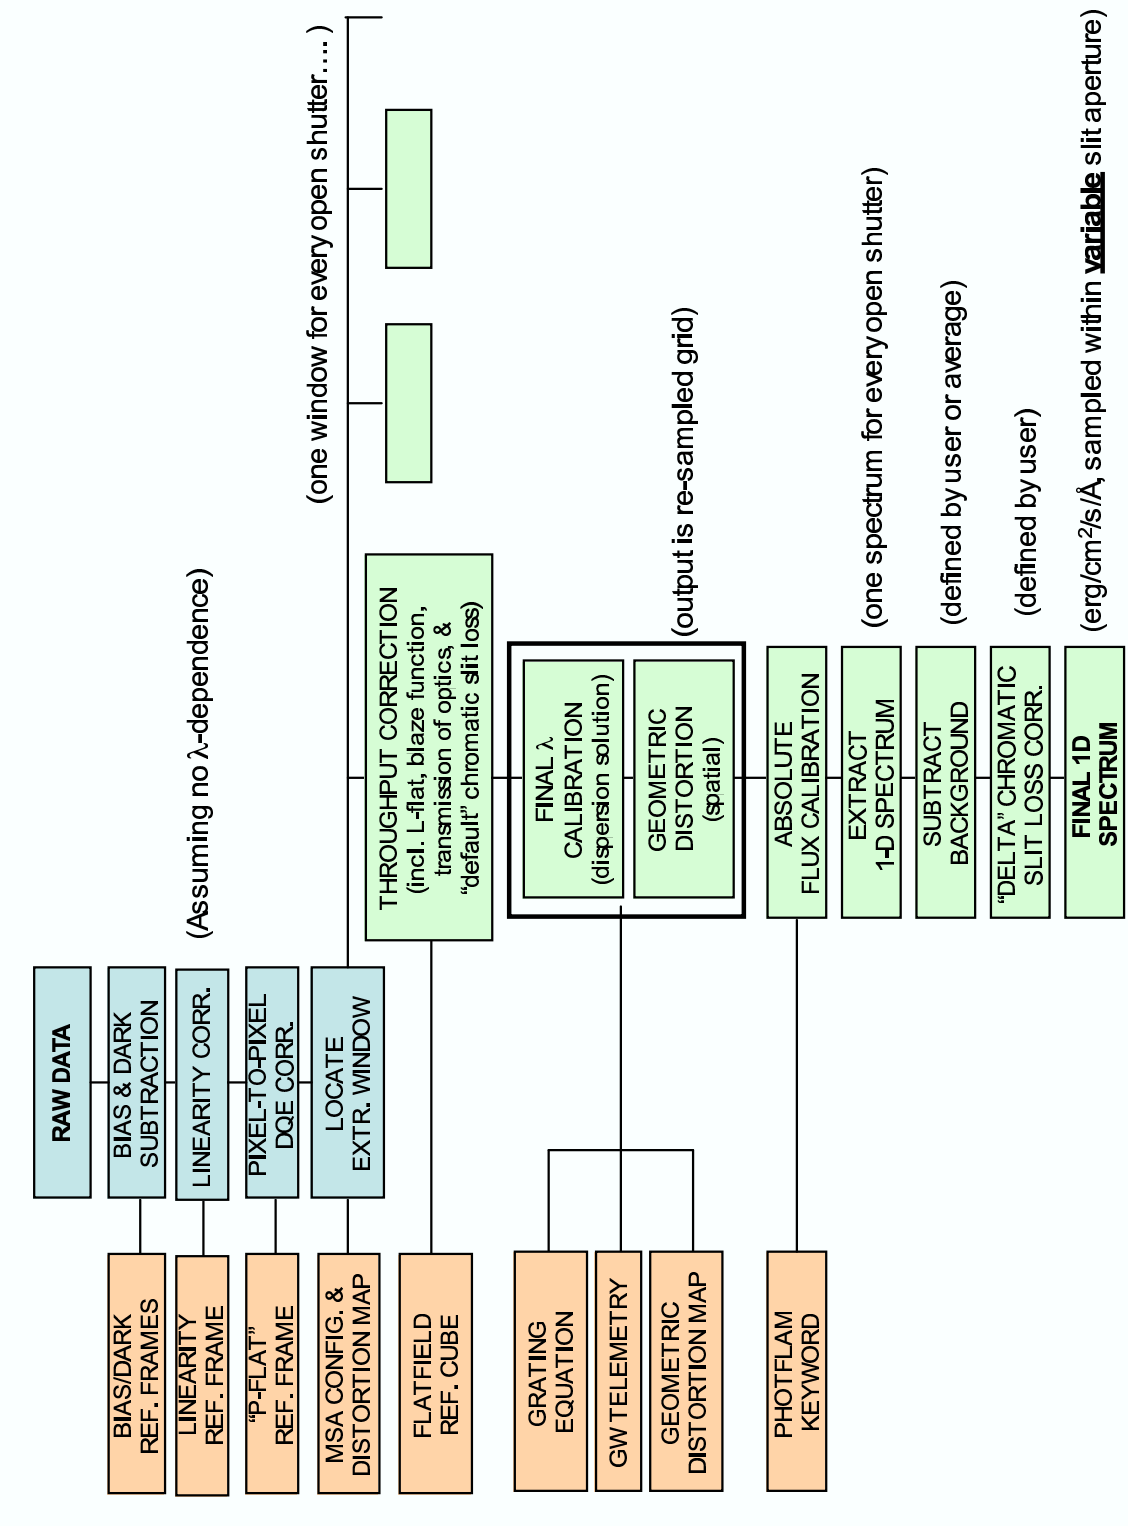

Figure 2: Schematic conceptual outline of the NIRSpec pipeline.

Part 5. Observatory and Two-Gyro Mode

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

HST Two-Gyro Mode

K. Sembach, M. Sirianni, S. Arribas, L. E. Bergeron, C. Biagetti, J. Biretta, G. Chapman, C. Cox, I. Dashevsky, R. de Jong, R. Doxsey, A. M. Koekemoer, V. Kozhurina-Platais, M. Lallo, R. Lucas, J. Mack, S. Malhotra, E. Nelan, K. Noll, C. Pavlovsky, C. Proffitt, M. Reinhart, K. Sahu, A. Schultz, A. Vick, T. Wiklind, & C. Xu

Space Telescope Science Institute, Baltimore, MD 21218

B. Clapp

Lockheed Martin Technical Operations Company., Greenbelt, MD 20770

Abstract. The Hubble Space Telescope is now operating with an attitude control system that relies upon inputs from two mechanical gas bearing gyroscopes and the Fine Guidance Sensors to provide fine pointing control during science observations. This mode of operation was developed to extend the operational lifetime of the observatory. Instrument performance in two-gyro mode is nominal, with no obvious degradation in performance compared to operations in three-gyro mode. In this review, we describe some recent instrument verification results and scheduling considerations for science operations in two-gyro mode.

1. Introduction to Two-Gyro Mode

The Hubble Space Telescope (HST) was originally designed to use three rate-sensing gyroscopes to provide fine pointing control of the observatory. In order to conserve the lifetime of the HST gyros, one of the functioning gyros was turned off on 28 August 2005, and a new attitude control system that functions with only two gyros was activated. In this mode, two gyros used in combination with the Fine Guidance Sensors provide fine-pointing information during science observations. The two gyros that are turned off could be reactivated should the need to do so arise (e.g., in the event of a failure of one of the two gyros currently in use).

On-orbit tests of the HST two-gyro fine guiding mode and its impact on science instrument performance were carried out on 20-23 February 2005 and 28-31 August 2005. More than 700 science exposures were obtained with the Advanced Camera for Surveys (ACS), the Near Infrared and Multi-Object Spectrograph (NICMOS), and the Fine Guidance Sensors (FGS) during the tests. All results from these tests indicate that there is no degradation in the quality of science data obtained in two-gyro mode compared to three-gyro mode.

In this article, we briefly describe the pointing jitter, instrument performance, and scheduling of observations in two-gyro mode. More information can be found in the *HST Two-Gyro Handbook* (Sembach et al. 2005) and in the Instrument Science Reports posted on the Two-Gyro Science Mode website at the following web address:

http://www.stsci.edu/hst/HST_overview/TwoGyroMode

Two-Gyro Mode Key Points:

- Science data obtained in two-gyro mode are essentially indistinguishable from those obtained in three-gyro mode.

- Observations requiring the finest pointing control (e.g., high-resolution imaging and coronagraphy) are feasible.
- The RMS fine-pointing jitter averaged over 60 seconds is typically ≤ 5 milli-arcseconds, which is much smaller than the detector pixels of the ACS and NICMOS.
- Scheduling is more restrictive in two-gyro mode because entry into fine-pointing mode for science observations is more complicated than in three-gyro mode.
- The following capabilities are not available in two-gyro mode:
  1. Gyro-only tracking
  2. Guide star handoffs
  3. Single guide star acquisitions
  4. Multiple roll positions within single, non-CVZ, orbits

2. Pointing Jitter

The HST Pointing and Control Systems Group measured the telescope pointing jitter during both the February 2005 and August 2005 on-orbit two-gyro tests using inputs derived from the gyros and FGS in the attitude control law to estimate the magnitude of the jitter. For each science exposure, the average and peak jitter found in 10-second and 60-second running intervals was calculated. Table 1 summarizes the mean and maximum values [in milli-arcseconds (mas)] found for each two-gyro dataset as well as the mean values for a set of three-gyro exposures obtained immediately prior to the February 2005 test.

Table 1: Pointing Jitter Summary

Mode	Gyro Set	# Exp.	Mean/Max RMS Jitter (mas)			
			10-sec Avg.	10-sec Peak	60-sec Avg.	60-sec Peak
Three-gyro[a]	1-2-4	24	4.1	5.2	4.2	4.3
Two-Gyro (Feb. 2005)	2-4	454	5.6 / 9.5	6.5 / 22.2	6.0 / 10.7	6.2 / 18.0
Two-Gyro (Aug. 2005)	1-2	262	3.3 / 5.3	3.9 / 11.7	3.4 / 4.9	3.6 / 5.3

[a] Three-gyro data were obtained 1 week prior to the February 2005 two-gyro test. Only mean values are listed.

The gyro set for each test was different, but the magnitudes of the jitter in both the two-gyro tests and in three-gyro mode are similar, as predicted by attitude control simulations prior to the tests. The pointing jitter in the current gyro configuration (gyros 1 and 2 + FGS) may be slightly less than it was in three-gyro mode. This difference is imperceptible in the highest resolution images taken with the ACS/HRC but is confirmed in data obtained with the FGS (see the contribution by E. Nelan et al., this volume). The predicted two-gyro jitter values for other combinations of gyros are similar to those reported in Table 1.

Increased jitter over short intervals resulting from impulsive spacecraft disturbances has been observed in both two-gyro and three-gyro mode. Most of these disturbances are so short (timescales of seconds) that the chances of seeing them in science data are small. Short science exposures have a low probability of occurring at the same time as the disturbances, and the jitter induced by the disturbances in long exposures contributes little to the overall S/N of the data. Descriptions of the types of disturbances and their frequency

of occurrence can be found in the *HST Two-Gyro Handbook* (Sembach et al. 2005). The present gyro configuration is insensitive to the common "V2-disturbances", which result from small rotations of the equipment shelf on which the gyros are mounted.

3. Instrument Performance

The Two-Gyro Science Mode Orbital Verification (TGSMOV) program in August 2005 is the most comprehensive set of tests of HST instrument performance in the current (Cycle 14) two-gyro configuration. Tests with both the ACS (programs 10458-10461) and NICMOS (programs 10462, 10464) provided information about the quality of high resolution imaging, pointing stability, coronagraphy, and moving target tracking (Table 2). The results of these tests are described below. A more complete description of the ACS PSF analysis and results can be found in ACS ISR 2005-11 (Sirianni et al. 2005). We refer the reader to the articles by Nelan et al. (this volume) for information about FGS astrometric measurements, and Koekemoer et al. (this volume) for information about pointing stability in two-gyro mode.

Table 2: August 2005 Two-Gyro Science Mode Orbital Verification Programs

Program[a]	Instrument	Purpose	Target
10458	ACS	PSF Shape/Stability/Dither	NGC 1851, NGC 2298, NGC 6752
10459	ACS	PSF Shape/Stability in CVZ	NGC 6752
10460	ACS	Coronagraphy	HD 216149
10461	ACS	Moving Target Tracking	Mars
10462	NICMOS	PSF Shape	NGC 1850
10464	NICMOS	Coronagraphy	HD 17925

[a]Program 10459 executed in October 2005. All others executed in August 2005.

The programs for the initial on-orbit verification of two-gyro fine pointing performance in February 2005 are listed in Table 3. Many of these tests served as the basis for the final TGSMOV programs listed above.

Table 3: February 2005 Two-Gyro Fine Pointing Test Programs

Program	Instrument	Purpose	Target
10443	ACS	PSF Shape/Stability/Dither	NGC 6341, Omega Cen
10444	ACS	PSF Shape/Stability in CVZ	CVZ-Field
10445	ACS	Coronagraphy	HD 130948A
10446	NICMOS	PSF Shape	P330-E
10447	NICMOS	Dither	NGC 6341
10448	NICMOS	Coronagraphy	GJ517

3.1. ACS Imaging

The ACS point spread function (PSF) shape and stability tests consisted of multiple exposures of three rich star clusters observed with the HRC and the F555W filter. Sequences of 10, 100, and 500 second exposures were obtained to check for dependencies of the PSF shape on exposure duration and timing within orbits. Observations using both V = 13 and V = 14 magnitude guide stars allowed for a check of the PSF width dependence on the

guide star magnitude. A total of 114 exposures for three clusters (NGC 1891, 2298, 6752) were obtained in August 2005 for orbits with occultations, and another 72 exposures for NGC 6752 were obtained in October 2005 while the cluster was in the continuous viewing zone (CVZ).

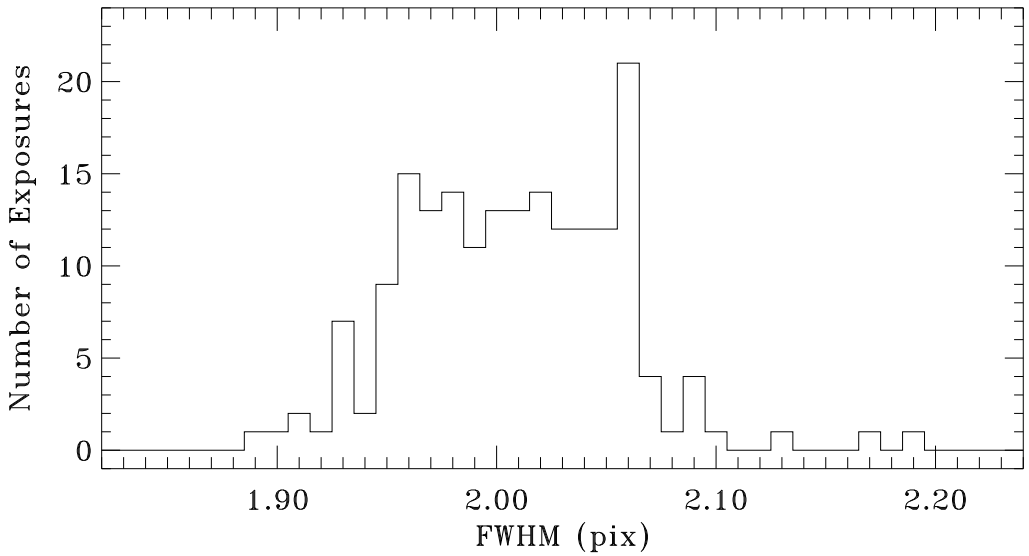

Figure 1: Histogram of point spread function widths in the ACS/HRC F555W filter during the August 2005 TGSMOV instrument performance test.

For each image, the FWHM of the PSF for stars with S/N > 10 was calculated by fitting Gaussian profiles to the stellar light profiles. This typically resulted in hundreds of measurements per image, which were then averaged to produce a mean PSF width for each image with an uncertainty of $\sim 0.05-0.15$ pixels, depending on the number and brightnesses of the stars used in the mean. Figure 1 shows the distribution of these averaged PSF FWHM measurements in units of HRC pixels. The HRC has a plate scale of $\sim 25''$ per pixel. The PSF FWHM values span a range of 1.89-2.19 pixels, with a mean of 2.00 pixels. These widths are comparable to the historical three-gyro average of 2.04 ± 0.03 pixels calculated in the same fashion.

The high-width outliers in Figure 1 are measurements for the cluster NGC 6752 taken in August 2005. These points can be compared to those on the low-width side of the histogram, which are measurements for NGC 6752 taken during October 2005. The August NGC 6752 widths range from 2.04 to 2.19 pixels with an average of 2.09 pixels, and the October widths range from 1.89 to 2.06 pixels with an average of 1.97 pixels. A possible reason for the difference in PSF widths for this cluster is the different Sun angles of the observations ($\sim 115^\circ$ in August versus $\sim 80^\circ$ in October). The observations for this cluster bound those seen in the two-gyro observations of the other clusters and set a practical estimate for the amount of scatter expected in measurements of the PSF width for a given pointing.

There were no apparent differences in PSF widths for exposures taken with bright (V = 13 mag) guide stars versus those obtained with faint (V = 14 mag) guide stars. However, there were differences in PSF widths seen for exposures taken at different times within orbits. The PSFs get broader with time within individual orbits suggesting that the dependence is likely due to normal changes in focus caused by the breathing cycle of the telescope during the orbit. Figure 2 shows a comparison of the FWHM measurements for a

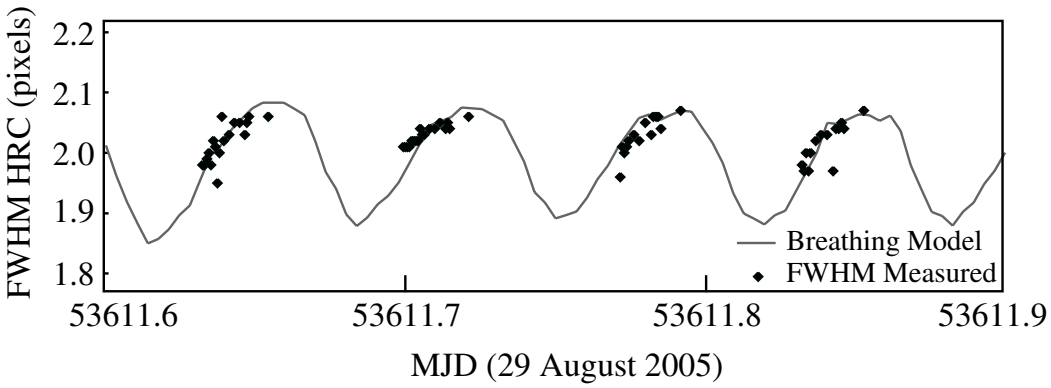

Figure 2: A model of the HST focus change induced by orbital telescope breathing compared to the point spread function widths measured for a portion of the ACS/HRC TGSMOV data.

portion of the TGSMOV data and a model of the expected changes caused by the breathing cycle. The typical PSF width change shown in this figure is ~ 0.1 pixel. The magnitude of this effect and the amplitude of the breathing model for any given pointing will depend upon the thermal stability of the observatory at that position.

Small changes in the PSF width due to exposure duration are also present in the TGSMOV data, indicating that longer exposures may have slightly larger PSF widths than shorter exposures. The measured widths for the August 2005 data were 2.02 ± 0.03, 2.05 ± 0.03, and 2.08 ± 0.05 pixels for the ensemble of 10, 100, and 500 second exposures, respectively. We note that the longer exposures were typically taken later in each orbit, so much of the difference is probably related to the intra-orbit focus changes mentioned above.

The same data used to check the PSF widths is suitable for checking the pointing stability within orbits. Figure 3 shows the exposure-to-exposure shifts in the pointing (V2, V3, total, and roll) as a function of time within the orbit for eight orbits in program 10458. These shifts were calculated by measuring the x and y positions of all stars in each image and comparing them to the positions of the stars in the first exposure within each orbit. The results of this test, which are summarized in Table 4, demonstrate that there is no significant difference in pointing stability between two-gyro and three-gyro mode. A detailed description of the procedure used to check the pointing stability and its application to the earlier February 2005 on-orbit test data can be found in ACS ISR 2005-07 (Koekemoer et al. 2005).

Table 4: ACS/HRC Pointing Stability

	Total Shift (RMS, mas)	Roll Angle (RMS, degrees)
Three-Gyro	2.19	0.00093
Two-Gyro (Feb. 2005)	2.29	0.00097
Two-Gyro (Aug. 2005)	2.08	0.00070

3.2. NICMOS Imaging

Even though the pixel sizes of the NICMOS cameras are larger than those of the ACS, data were obtained in TGSMOV program 10462 to check both the point spread function shape

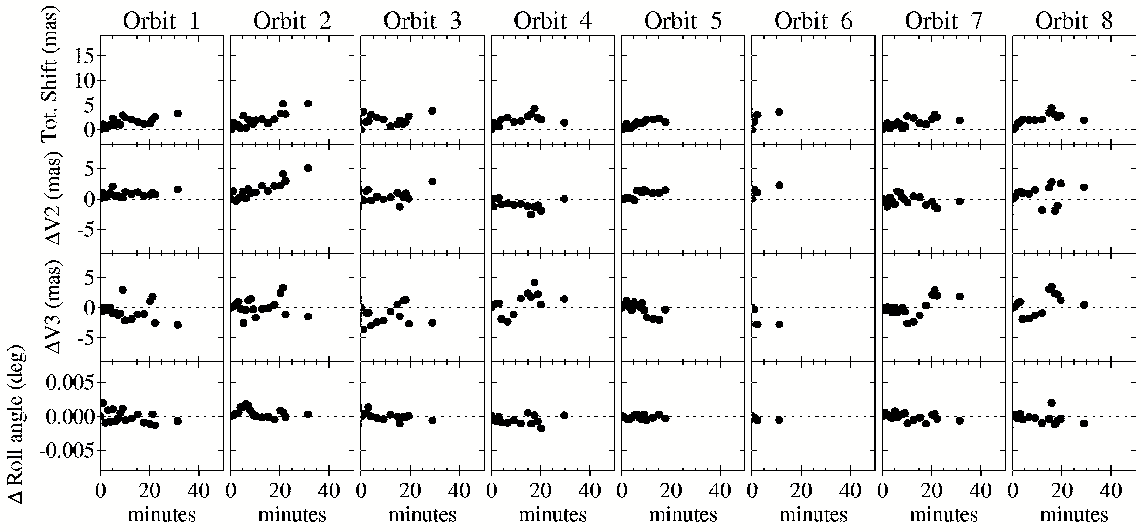

Figure 3: ACS/HRC pointing stability within 8 orbits of TGSMOV program 10458. Shifts are measured with respect to the first exposure in each orbit.

and dither pattern positioning. Measurements of the PSF in the NIC1 channel show no significant differences with those taken previously in the February 2005 test or in three-gyro mode. Dither pattern positioning and repetition in two-gyro mode are also as reliable as in three-gyro mode.

3.3. Coronagraphy

Several tests in February 2005 and August 2005 verified the ability to perform ACS and NICMOS coronagraphic acquisitions in two-gyro mode. The procedures for performing these acquisitions were the same as those used previously in three-gyro mode. Examples of coronagraphic images for ACS are shown in Figure 4, which compares the three-gyro and two-gyro observations of the star HD 130948A with the two-gyro observation of HD 216149. Figure 5 shows a comparison of the NICMOS F110W and F160W direct and coronagraphic images of HD 17925. In both the ACS and NICMOS data, the star light is nulled to the same extent as in three-gyro mode. Quantitative information about the coronagraphic nulling as a function of position from the center of the coronagraphic spot for the February 2005 two-gyro test can be found in NICMOS ISR 2005-001 by Schneider et al. (2005) and ACS ISR 2005-05 by Cox & Biretta (2005).

A primary limitation in two-gyro mode is the inability to perform observations at more than one roll angle within a single occulted orbit. The primary reason for this limitation is that the attitude control system must revert to a coarser tracking mode during the roll maneuver (unlike three-gyro mode), and there simply is not enough time to perform two complete sets of guide star and target acquisitions within a single orbit.

3.4. Moving Target Tracking

The TGSMOV observations for program 10461 verified the ability to track moving targets in two-gyro mode. The test consisted of 32 0.3-second exposures of Mars taken over the course of an orbit. Shifts were calculated for each exposure by cross-correlating the edges of the planetary disk. The rotation of Mars moved a number of bright surface features across the limb and the terminator during the course of the observations. The planetary rotation and gradual change in brightness of the day-night terminator complicated identification of the disk edge and limited the accuracy of the resulting cross-correlations. Comparison of the

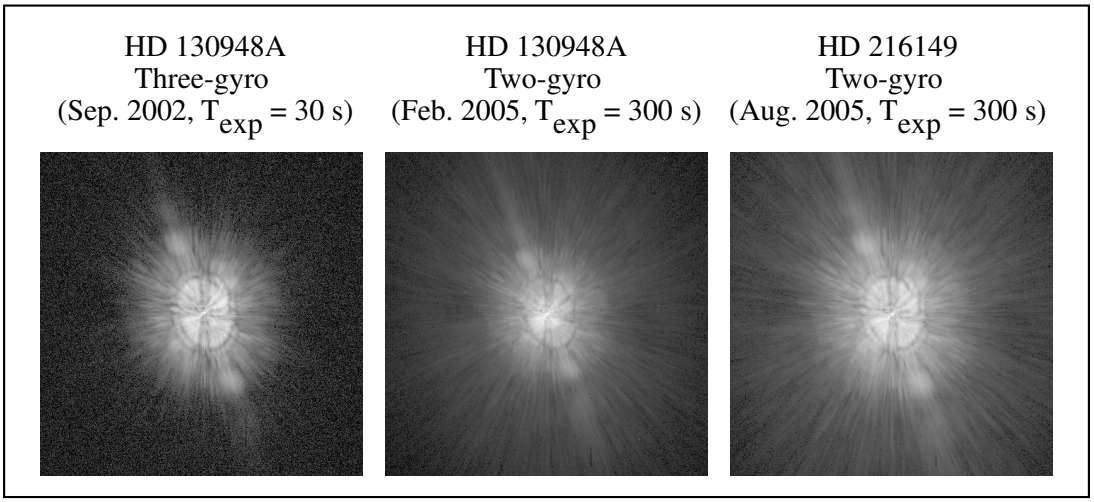

Figure 4: ACS coronagraphic images in three-gyro mode (*left*) and two-gyro mode (*center* and *right*). The three-gyro image of HD 130948A (30 sec) is grainier than the two-gyro image (300 sec) because the color scaling is stretched to reveal weaker features that are more difficult to see in this shorter exposure.

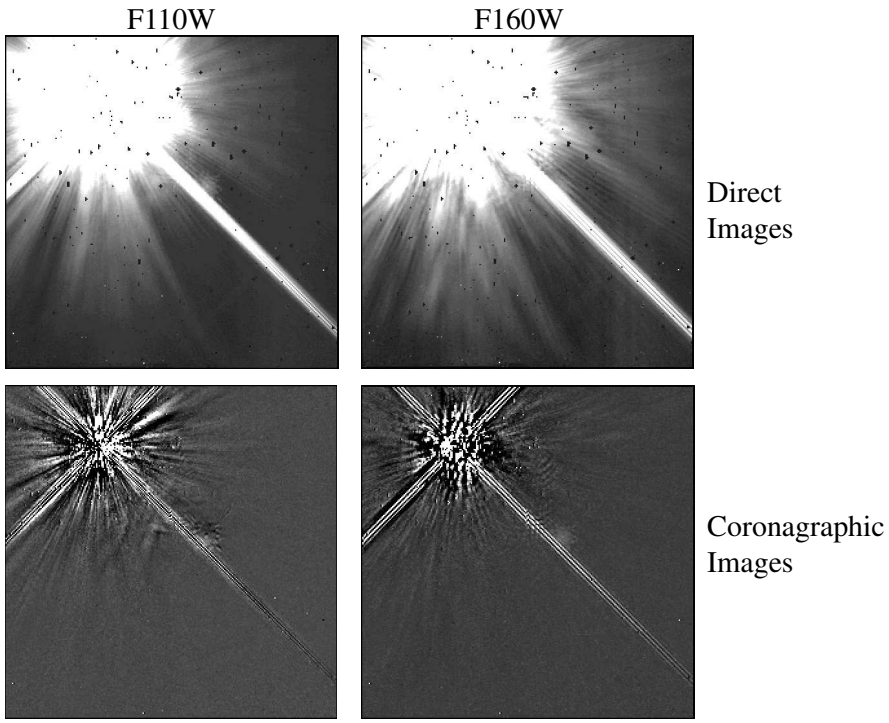

Figure 5: NICMOS F110W and F160W direct and coronagraphic images of HD 17925 in two-gyro mode (August 2005).

measured shifts between exposures with those expected from the predicted and final HST ephemeris yielded residuals smaller than the unavoidable errors resulting from in-track HST positional uncertainties. Therefore, while not specifically a test of instrument performance,

this test verified that it is possible to track observable solar system objects with a precision comparable to that in three-gyro mode.

4. Scheduling of Observations

Observations with either orientation or timing constraints are more difficult to schedule in two-gyro mode than in three-gyro mode because of the additional pointing restrictions necessary for attitude control and observatory safety. Roughly half the sky is visible at any point in time in two-gyro mode, compared to $> 80\%$ of the sky in three-gyro mode. For unconstrained observations (those having neither a timing constraint nor an orientation constraint), every point in the sky is observable at some time during the year.

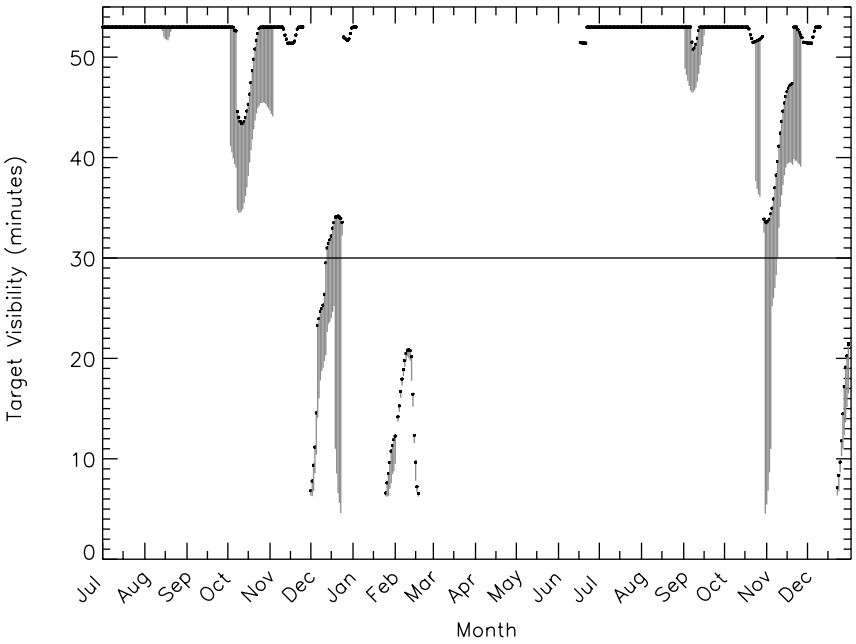

Figure 6: Two-gyro visibility for the Hubble Ultra Deep Field in Cycle 15 (beginning July 2006). Black points indicate the maximum orbital visibility, and gray lines indicate the range of visibilities for allowed observatory roll angles.

Constrained observations have visibilities that depend upon the time of year and the declination of the object to be observed. An example of the visibility for the Hubble Ultra-Deep Field at $\alpha = 3^h32^m$, $\delta = -27°55^m$ is shown in Figure 6. Tools that provide detailed visibility information for every point on the sky are available at the HST Two-Gyro Science Mode website. Descriptions and examples of how orientation and timing constraints can affect the scheduling of an object and its orbital visibility period can be found in the *HST Two-Gyro Handbook*.

Scheduling efficiency in two-gyro mode is slightly lower than in three-gyro mode, with approximately 70 prime science orbits per week being scheduled compared to 80 per week in three-gyro operations. These numbers do not include orbits used for snapshots or calibration programs. It may be possible to improve the two-gyro scheduling efficiency by relaxing some of the scheduling constraints once more experience is gained with two-gyro operations.

5. Guide Star Acquisition Success Rate

Guide star acquisitions in two-gyro mode are presently slightly more susceptible to failure than in three-gyro mode because the attitude control system may revert to a coarser pointing mode if problems arise during the acquisition process. Initial indications from the first few months of two-gyro operations are that the guide star acquisition success rate in two-gyro mode in the September to November 2005 timeframe ($\sim$ 97%) is nearly as high as the historical three-gyro mode average ($\sim$ 99%). A variety of factors may contribute to the slight difference in acquisition success rates. First, the previous three-gyro success rate was achieved by finely tuning the attitude control law over many years, whereas the two-gyro success rate is based upon a limited time sample. A longer period of time is needed to track progress in two-gyro mode. Second, many of the two-gyro acquisition failures to date have been due to problems associated with the pointing mode used prior to entering fine lock. The prior pointing mode uses two gyros and the fixed-head star trackers to perform the onboard attitude determination before the FGSs can be used in the pointing control loop. Those problems are being studied and are in the process of being corrected in the flight software. Other failures related to the acquisition logic are also being revised. Once these issues are corrected, the two-gyro acquisition success rate should converge to the historical three-gyro rate because the only remaining failures should be the usual ones related to scheduling "bad" guide stars that the FGS probably has no chance of acquiring in either two-gyro or three-gyro mode.

Acknowledgments. The development and implementation of the HST Two-Gyro Mode was a team effort involving many people - too many to name here. We are indebted to all of our colleagues at STScI, GSFC, and elsewhere who worked so hard to make this mode of operation possible. We are also grateful to the HST Pointing and Control Systems Group for their willingness to share their pointing jitter results.

References

Cox, C. & Biretta, J. 2005, *ACS Coronagraph Performance in Two-Gyro Mode*, *Instrument Science Report* ACS 2005-05, (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Koekemoer, A., Kozhurina-Platais, V., Riess, A., Sirianni, M., Biretta, J., & Pavlovsky, C. 2005, *Two-Gyro Pointing Stability of HST Measured with ACS*, *Instrument Science Report* ACS 2005-07, (Baltimore: STScI)

Schneider, G.A., Schultz, A., Malhotra, S., & Dashevsky, I. 2005, *NICMOS Two-Gyro Mode Coronagraphic Performance*, *Instrument Science Report* NICMOS 2005-001, (Baltimore: STScI), available through http://www.stsci.edu/hst/nicmos

Sembach, K.R., et al. 2005, *HST Two-Gyro Handbook*, Version 2.0, (Baltimore: STScI)

Sirianni, M., et al. 2005, *Characterization of the ACS/HRC Point Spread Function in Two-Gyro Mode*, *Instrument Science Report* ACS 2005-11, (Baltimore: STScI)

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The Two-Gyro Pointing Stability of HST, Measured with ACS

A. M. Koekemoer, V. Kozhurina-Platais, M. Sirianni, A. Riess, J. Biretta and C. Pavlovsky

Space Telescope Science Institute, 3700 San Martin Dr., Baltimore MD 21218, USA

Abstract. We present the results of the pointing stability tests for HST, as measured with the ACS/HRC during the two-gyro test program conducted in February 2005 and the validation tests performed after the transition to two-gyro mode in August 2005. We have measured the shifts of 301 exposures obtained of the globular clusters NGC1851, NGC2298, NGC6341, NGC6752, and Omega Centauri, obtained over a total of 21 orbits during these two test programs, and compare the measured pointings to those that were commanded in the observing program. We find in all cases that the measured shifts and rotations have the same level of accuracy as those executed in three-gyro mode. Specifically, the pointing offsets during an orbit relative to the first exposure can be characterized with distributions having a dispersion of 2.1 - 2.3 milliarcseconds for shifts and 0.0007 - 0.00097 degrees for rotations, thus less than 0.1 HRC pixels, and agree extremely well with similar values measured for comparable exposures obtained in three-gyro mode. In addition, we successfully processed these two-gyro test data through the MultiDrizzle software which is used in the HST pipeline to perform automated registration, cosmic ray rejection and image combination for multiple exposure sequences, and we find good agreement with similar exposures obtained in three-gyro mode. In summary, we find no significant difference between the quality of HST pointing as measured from these two-gyro test data, relative to the nominal behavior of HST in regular three-gyro operations.

1. Introduction

As part of the options for extending the lifetime of the Hubble Space Telescope (HST), a new attitude control system was activated in August 2005 that enables the telescope to point using only two gyroscopes instead of the nominal three by using additional information from the Fine Guidance Sensors, thereby allowing one of the gyroscopes to be turned off to conserve its life for possible future use. The February two-gyro test (F2G) was carried out during $20 - 23$ February 2005 to investigate the feasibility of this mode, and based on its results the decision was made to transition to two-gyro mode (TGM) on 28 August 2005, with a second set of validation tests performed during $28 - 31$ August 2005.

One of the principal motivations for these tests was to examine the pointing stability of HST, since a possible concern with TGM operations is the need to determine whether the pointing stability is significantly worse than with three gyroscopes. If so, this could potentially impact the tracking accuracy of HST during an exposure, as well as the accuracy of dither offset maneuvers from one exposure to the next, and the accuracy of guide star re-acquisitions from one orbit to the next. In addition to degrading the resolution of the final image, a decrease in pointing stability can also affect the degree of accuracy with which sub-pixel sampling is achieved with dithering. Therefore, an extensive set of observations was obtained during both test periods to verify and measure the pointing stability of HST under TGM. All results from these tests indicate that the two-gyro pointing performance of HST is not significantly different from that under nominal three-gyro operations.

2. Observational Design and Methodology

The observational methodology is described in detail in ACS ISR 2005-07 (Koekemoer et al. 2005) and Sembach et al. (this volume), and is summarized here. Three principal questions need to be addressed concerning the pointing stability of HST in two-gyro mode:

1. Maintaining stability within an orbit, so that successive exposures during an orbit are either located at the same position (if no dithering is used), or are dithered accurately according to the commanded offsets, particularly if sub-pixel shifts are required.

2. Maintaining sufficient re-acquisition accuracy from one orbit to the next to enable multi-orbit observations to be successfully obtained, particularly if sequences of exposures or dither patterns are spread across multiple orbits.

3. Ensuring that the pipeline processing system, specifically MultiDrizzle image registration, cosmic ray rejection and image combination, can successfully process the data.

The observations described here were obtained in programs 10443 and 10458 during February and August 2005, respectively. These contained a wide variety of exposures, using a range of exposure times, dither offset strategies and different types of guide stars, to quantify in detail the behavior of the telescope in TGM. The proposals used the HRC camera (1024x1024 pixels), which has small pixels (26 milliarcseconds) and can very accurately measure the HST pointing accuracy. The 10443 observations were spread over 13 orbits, obtaining 155 exposures of the globular cluster NGC6341 and 32 exposures of Omega Centauri at a 10° off-nominal roll. The 10458 program covered 8 orbits and obtained 18 exposures of NGC1851, 18 exposures of NGC6752, and 78 exposures of NGC2298. A companion three-gyro program, 10455, was obtained in February 2005 using a subset of the exposures in 10443 with the same observing configuration, therefore providing useful exposures that could be directly compared with those obtained in two-gyro mode.

The observations used a range of guide stars with magnitudes V=11, 13, and 14. The exposure times were 10 seconds, 100 seconds and 500 seconds, and were obtained in a variety of configurations including CRSPLIT sequences of 2, 4 and 5 exposures (with no dithering) as well as 2-point and 4-point dither patterns, with some of the 2-point patterns containing a 2-exposure CRSPLIT pair at each location. Most of the observations used the F555W filter, with additional F330W observations obtained in February 2005.

3. Analysis and Results

3.1. Initial Processing and Distortion Correction

All the exposures were first processed through standard ACS calibration, including gain correction, bias and dark current removal, and flat field correction. The resulting calibrated FLT files were then transformed onto an undistorted output frame using the MultiDrizzle software (Koekemoer et al. 2002), which makes use of Drizzle (Fruchter & Hook 2002) to remove the ACS geometric distortion using the most up-to-date distortion files (IDCTAB) and distortion residual images (DGEOFILES), as specified in the image headers (ACS ISR 2004-15, Anderson & King 2005). This step also accounts for slightly different distortion terms in different filters (F555W and F330W), and removes small additional scale changes due to velocity aberration resulting from changes in the motion of HST along the line of sight to the target during an orbit. Thus, exposures from different times and with different filters could be directly compared with one another. The resulting set of drizzled images, one for each exposure, were all examined in detail to verify that there were no problems, before continuing with the pointing measurements.

3.2. Catalog Generation

The pointing accuracy was measured for all the exposures, regardless of whether they had been obtained as part of a CRSPLIT or NUMEXP sequence with no dithering, or whether a dither pattern had been used. The goal was to measure how well each exposure aligned with the commanded pointing of the telescope. This was measured by first creating a catalog for each exposure using the IRAF DAOFIND software (Stetson 1987), with parameters optimized for centroiding on unresolved stellar sources. For the clusters NGC2298, NGC1851, NGC6341, and NGC6752, the images typically contained $\sim 1000 - 2,500$ stars that could be matched between all the exposures (10, 100 and 500 seconds) while the F330W images contained $\sim 800 - 1,100$ stars that could be matched. For Omega Centauri, only $\sim 50 - 60$ stars were matched so this dataset served more as a consistency check. Cosmic rays were generally not a problem; most of the exposures were short enough to have a low number of cosmic rays, and the number of stars was generally large enough that occasional stars affected by cosmic rays would show up as significant outliers and could be easily rejected.

3.3. Shift Measurement Results

We used the catalogs to iteratively solve for shifts, rotations and possible scale changes. All scale changes from velocity aberration were successfully removed, and no significant additional time-dependent scale changes were present. The only remaining scale change is between short and long exposures, along the y-axis to the level of a few times 10^{-5}. This is accounted for by a known effect related to CTE (charge transfer efficiency) which produces slight changes in the centroids of stars. However, this scale change is not time-dependent, and comparisons between exposures of the same length showed no significant scale changes during the observations. No other geometric changes were found and the final analysis was conducted by solving for shifts and rotations, keeping other terms fixed. The resulting measurements for each exposure represent the difference between the commanded HST pointing and the actual pointing obtained. All of these can be presented relative to the first exposure in each orbit, and are shown in Figure 1.

The results show that the stability of the telescope is generally very good, with each exposure aligning to within a few milliarcseconds of the commanded offsets. While some orbits show a slight gradual change with time up to $\sim 4 - 5$ milliarcseconds, this is also found to occur in three-gyro mode and is interpreted as the result of thermal changes within the telescope. It should be recalled that the orbits during each test were not contiguous but obtained over several days, during which other targets at different sun angles were also observed, therefore slight changes in tracking due to thermal effects might be expected. However, the principal result is that the translational and rotational stability of the telescope remains good and shows no dramatically different behavior to three-gyro mode.

We also compared these results to three-gyro data from program 10455 and Cycle 12 program 9750 (PI: K. Sahu), which used ACS/WFC over 105 orbits to observe the galactic bulge (TEL ISR 2005-02, Gilliland et al. 2005). Since the telescope becomes thermally stable after about a day, we chose the first 20 orbits of this program as being representative of three-gyro data. In Figure 2 we show a histogram summarizing the relative offsets that were presented in Figure 1 for two-gyro mode, along with the similar measurements obtained for the comparable data from program 9750. In all cases, the shifts and rotations represent the differences from the commanded pointings for all exposures within an orbit, relative to the first exposure in the orbit. From the figure it is apparent that the distributions are not significantly different, and this is also borne out by the quantitative comparison between the two distributions when we characterize each distributions in terms of its r.m.s. dispersion. This was verified independently by A. Riess and the other team members. The results are presented in Table 1. We also find that the pointing repeatability for guide star re-acquisitions between orbits does not appear significantly worse than in three-gyro mode, although we were only able to verify this for a few orbits.

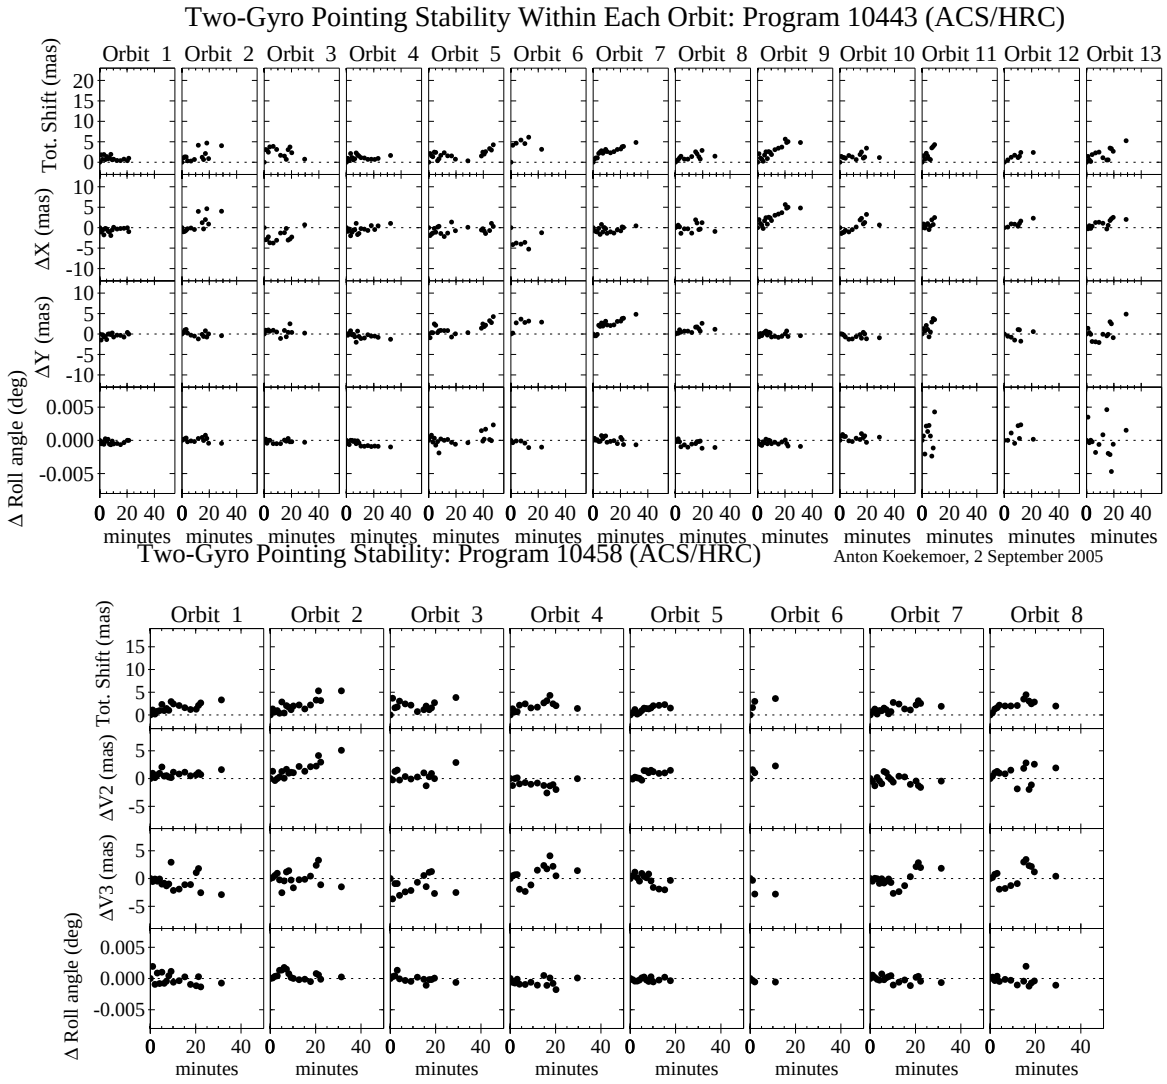

Figure 1: Relative pointing accuracy of HST for the observations obtained in February 2005 (top) and August 2005 (bottom), for a total of 301 exposures, including dither patterns. The offsets represent the difference between the commanded and actual pointing, to quantify the stability of HST during an orbit. It can be seen that the stability is very good to the level of a few milliarcseconds, which is comparable to three-gyro performance.

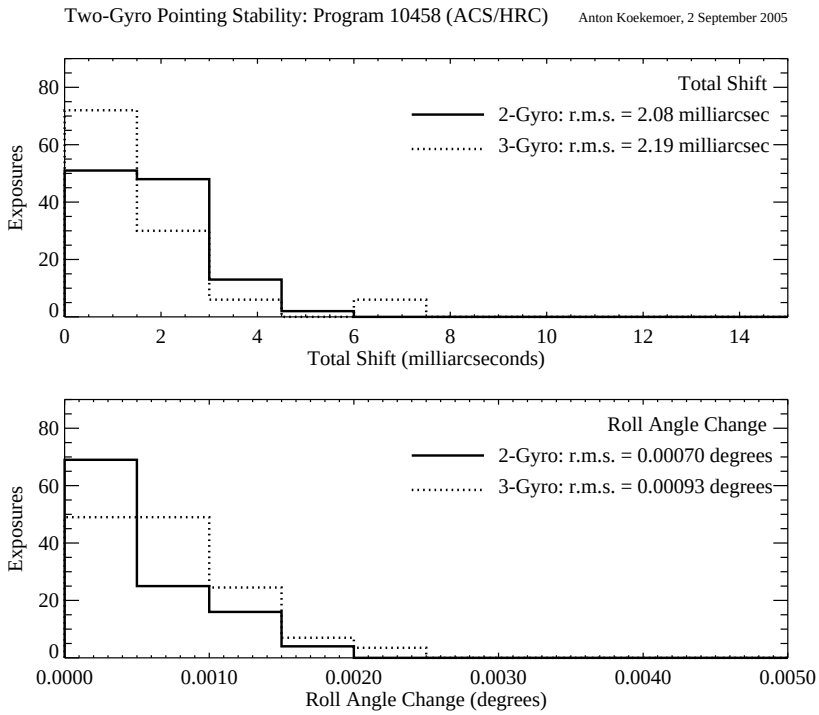

Figure 2: A comparison of the histogram distribution of shifts and rotational offsets between the two-gyro test data and comparable three-gyro data. These offsets represent the differences between commanded and measured pointings, as plotted in Figure 1. There is no significant difference between two-gyro and three-gyro data for either shifts or rotations.

Table 1: HST Pointing Stability Summary: ACS/HRC

	Total Shift r.m.s. (milliarcseconds)	Roll Angle r.m.s. (degrees)
Two-gyro (February 2005)	2.29	0.00097
Two-gyro (August 2005)	2.08	0.00070
Three-gyro (nominal)	2.19	0.00093

3.4. Processing Two-Gyro Data through MultiDrizzle

As a final test, we processed the two-gyro exposure sequences through the MultiDrizzle software (Koekemoer et al. 2002), which is used in the HST pipeline to automatically register images based on their headers, perform cosmic ray cleaning and create a final combined image using Drizzle. A variety of multiple exposure sequences were obtained in the two-gyro test, including those specified using CRSPLIT, NUMEXP, and dither patterns. As an initial test, the images were processed through MultiDrizzle using only their header astrometry as a basis for registration, thus ignoring the offsets of a few milliarcseconds that were shown in Figure 1. This is the current behavior of the pipeline, in the sense that the images are registered based on their commanded offsets. The results from these tests showed that there is no significant degradation in the quality between the images, which was confirmed quantitatively by measuring the PSF of stars in the images in each case.

This agreement is consistent with the fact that the offsets of the exposures as shown in Figure 1 were found to be only a few milliarcseconds, or less than about 0.1 HRC

pixels. Therefore, these two-gyro images can generally be combined directly to the same degree of accuracy as three-gyro data, using the astrometric headers information. If it is ever necessary to incorporate measured shifts, as demanded by certain types of scientific programs, then this can be done by means of the "Tweakshifts" script which uses the techniques described here to solve for shifts and apply them to the image headers, prior to image combination.

4. Summary

We have presented the results of the pointing stability tests for HST, as measured with the ACS/HRC during the Two-Gyro test program conducted in February 2005 as well as the verification observations in August 2005. We have measured the shifts of 301 exposures of the globular clusters NGC1851, NGC2298, NGC6341, NGC6752, and Omega Centauri, obtained over a total of 21 orbits, and have compared the measured pointings to those that were commanded in the observing program. We find in all cases that the measured offsets of shifts and rotations agree with those that were commanded to the same level of accuracy as in three-gyro mode. Specifically, the differences between commanded and actual pointings during an orbit relative to the first exposure can be characterized with distributions having dispersions of $\sim 2.1 - 2.3$ milliarcseconds for shifts and $\sim 7.0 - 9.7 \times 10^{-4}$ degrees for rotations, thus less than 0.1 HRC pixels, and agree extremely well with similar values measured for comparable exposures obtained in three-gyro mode. In addition, we successfully processed these two-gyro test data through the MultiDrizzle software which is used in the HST pipeline to perform automated registration, cosmic ray rejection and image combination for multiple exposure sequences, and we find excellent agreement with similar exposures obtained in three-gyro mode. In summary, we find no significant difference between the quality of HST pointing as measured from these two-gyro test data, relative to the nominal behavior of HST in regular three-gyro operations.

Acknowledgments. We are grateful to the ACS group and the operations teams at STScI and GSFC who worked to implement two-gyro operation, and to Ron Gilliland for kindly sharing the data from his ISR in order to facilitate a comparison with three-gyro observations.

References

Anderson, J., & King, I., 2004, "Multi-filter PSFs and Distortion Corrections for the HRC", *Instrument Science Report* ACS 2004-15 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Fruchter, A. S. & Hook, R. N., 2002, PASP, 114, 144

Gilliland, R. 2005, "Guiding Errors in 3-Gyro: Experience from WF/PC, WFPC2, STIS, NICMOS, and ACS", *Instrument Science Report* TEL 2005-02 (Baltimore: STScI)

Koekemoer, A. M., Fruchter, A. S., Hook, R. N., & Hack, W.,2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 337

Koekemoer, A. M., Kozhurina-Platais, V., Riess, A., Sirianni, M., Biretta, J., & Pavlovsky, C. 2005, "Two-Gyro Pointing Stability of HST Measured with ACS", *Instrument Science Report* ACS 2005-07 (Baltimore: STScI)

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

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

FGS Astrometry in Two-Gyro Mode

E. Nelan

Space Telescope Science Institute

Abstract. The Fine Guidance Sensor 1R on board *HST* is used by observers as a science instrument for state of the art relative astrometry and for the high angular resolution of close binary systems. The post observation analysis of FGS science data requires the removal, among other things, of spacecraft jitter and drift that occurred during the observations. Under three-gyro mode this technique had been perfected and implemented as part of the FGS astrometry calibration pipeline. Special tests were conducted in February 2005 to assess FGS astrometric performance in two-gyro mode. No degradation was noted, and no modifications to the calibration pipeline were found to be necessary. FGS science and calibration data acquired since the transition to two-gyro operations in August 2005 confirm the results of the February test. However, in two-gyro mode the scheduling of observations is more restrictive. This especially impacts parallax programs of distant targets at low to moderate declinations since such fields can be observed at only one of the two epochs of maximum parallax factor.

1. Introduction

The *HST* Fine Guidance Sensor 1r (FGS1r) interferometer has been calibrated as an astrometric science instrument. In this capacity it has two modes, Position and Transfer. Position mode is used for wide angle (up to 69 deg^2) relative astrometry whereby the instrument sequentially acquires and tracks the fringes of selected stars residing in its field of view (FOV). In Transfer mode, which can be used to resolve small scale structure, e.g., close binary systems, FGS1r repeatedly scans its 5"x5" instantaneous FOV across a target to obtain data from which the object's interference fringes can be reconstructed (refer to the *FGS Instrument Handbook* for details: Nelan 2005). The post observation processing of both Position and Transfer mode data require the removal of spacecraft jitter and drift that occurred during the course of the observations, which typically span the full 55 minutes of target visibility during an *HST* orbit. The removal of jitter is facilitated by correlating the guide star centroids provided by the guiding FGSs with the centroids from the measurements of the astrometric targets. In Position mode, the measured positions stars which have been observed multiple times over the course of the visit are used to model drift as a second order time dependent polynomial. This model is reverse-applied to the centroids of all the target stars (which removes the affect to the observatory drift). In Transfer mode, drift is removed by cross correlating the observed fringes from the individual scans. Jitter and drift removal is an automated and routine part of FGS science calibration pipeline that was perfected under three-gyro operations.

During the development of the two-gyro mode capability, it was generally expected that spacecraft pointing performance under two-gyro mode would be somewhat degraded relative to three-gyro operations, especially along the axis of the "missing" gyro. However, this did not necessarily imply that the quality of FGS astrometry would be degraded since we expected the jitter and drift corrections of the calibration pipeline to be robust to the new spacecraft pointing characteristics. To verify this, 4 orbits of FGS1r astrometry

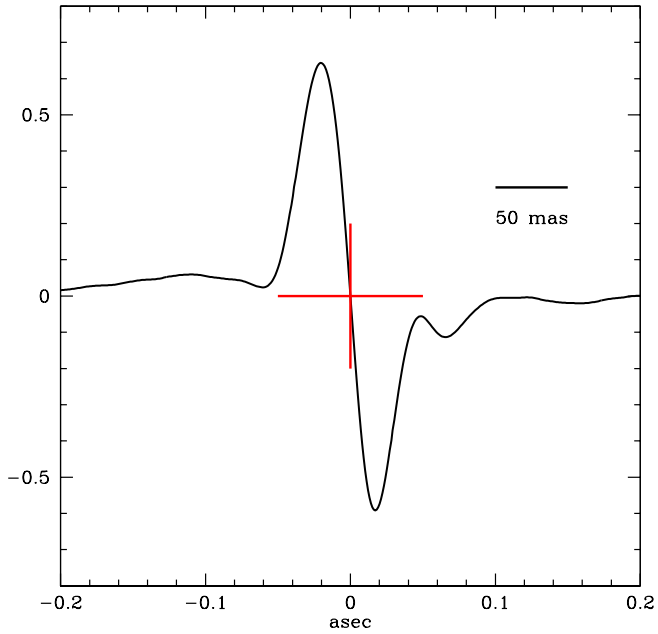

Figure 1: The FGS interference fringe. Vertical direction is fringe amplitude, horizontal direction is angle on the sky. Note the indicated angular scale. During a Position mode observation or while tracking a guide star, the FGS tracks the feature of the fringe at the cross hairs.

were included as part of the February 2005 on-orbit two-gyro mode test (Nelan 2005). Since the transition in late August 2005 to full time two-gryo operations, additional FGS astrometry data from both the GO and calibration programs have become available. The recent data confirm the findings of the February test and shows that the calibration pipeline is indeed well suited for processing two-gyro mode astrometry data. Interestingly, we find that contrary to earlier expectations, there is no significant degradation to the quality of the *HST* pointing control system performance relative to three-gyro operations. We expand upon this in the next section.

2. FGS Astrometry Data in Two-Gyro Mode

In Position mode the FGS tracks the so called "zero-point" crossing of the star's interference fringe (Figure 1). As the position of the star in the FGS detector frame changes, the star selector servos reposition the instrument's optical axis to keep the "cross hair" centered on the fringe. From these adjustments, which occur at 40 Hz (for bright stars), one can track the 2-dimensional (x,y) position of the star in the detector during the course of the observation (the FGS generates fringes in two orthogonal directions).

The guiding FGSs track their respective guide stars in an identical fashion, an important difference being that the guide star centroids are used by the spacecraft pointing control system to fine point and stabilize *HST*. By converting the local (x,y) coordinates of the astrometry target and guide star centroids in the respective FGSs into the vehicle (V2,V3) reference frame, one can compare how well changes in the centroids from the 3 FGSs correlate. A close correlation indicates that each FGS observed changes corresponding to motion of the telescope, while poor correlation indicates that the *noise equivalent angle* of one or more of the centroids is larger than the pointing jitter.

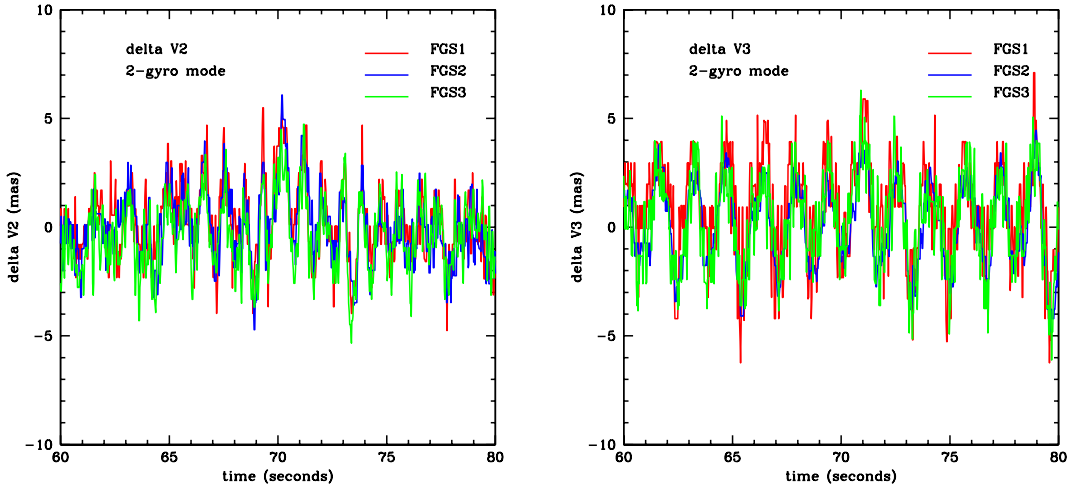

Figure 2: The correlation of V2 (left) and V3 (right) jitter witnessed by the 3 FGSs in a two-gyro mode observation.

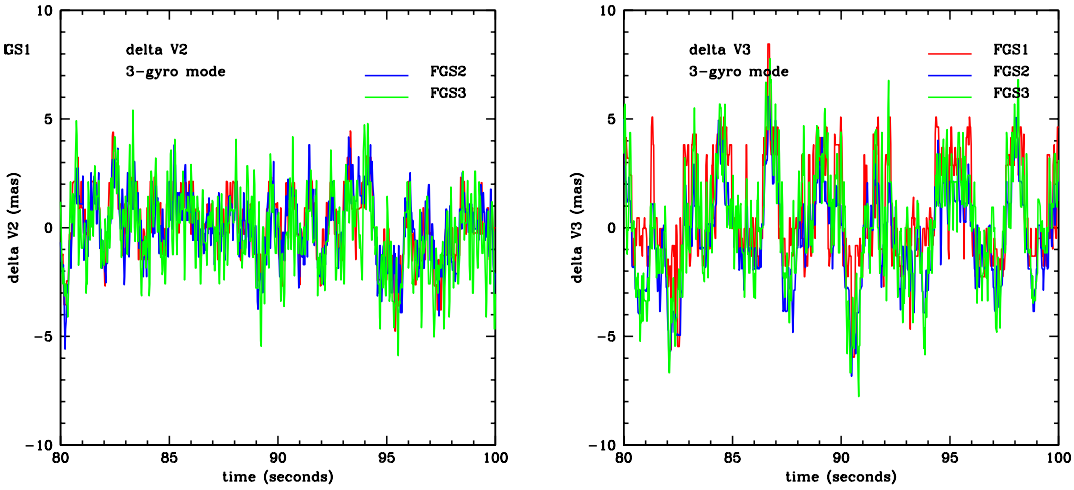

Figure 3: The correlation of V2 (left) and V3 (right) jitter witnessed by the 3 FGSs in a three-gyro mode observation.

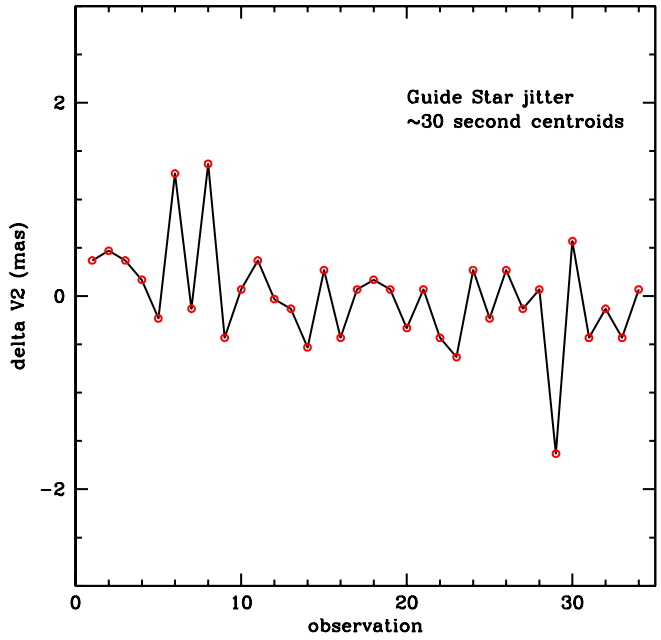

Figure 4: The change in the V2 centroid of the dominant guide star during the course of an astrometry visit in two-gyro operations. The circles mark the temporal midpoints of the individual astrometry exposures. At each point, the guide star centroids are computed over the time of the astrometry exposures. The stability of the guide star centroids in two-gyro mode is essentially that same as in three-gyro mode. This "jitter" is removed by the calibration pipeline.

2.1. Jitter

Figures 2 & 3 show the 40 Hz correlation of an astrometry target star and the two guide star centroids in V2 and V3 for observations under two-gyro and three-gyro operations, respectively. In these two cases all three stars are fairly bright and the data from the three FGSs correlate quiet well. It is note worthy to point out that the amplitude of the jitter is similar in both cases. The apparent periodicity of the jitter is due a vibrational mode of the high gain antennae and is present in both two and three gyro operations. The good correlation across the three FGSs indicates that the guidance data can be used to eliminate jitter in the astrometry data.

However, the calibration pipeline does not use the 40 Hz guidance data for jitter removal. Rather, the pipeline uses the guide star centroids that are computed (via a trimmed mean or a median filter) over the duration of an individual astrometry observation, which is typically 20 to 30 seconds. For illustration, Figure 4 shows the change in the V2 location of a guide star's position in the dominant guider (FGS2 in this case, corrected for differential velocity aberration) over the course of an orbit while FGS1r executed astrometry observations. The small circles mark the mid-points of the astrometry observations. The data plotted here are from two-gyro operations. This is consistent with three-gyro performance.

2.2. Drift

The drift of the *HST* focal plane across the sky during the course of a visit is a familiar phenomenon that is generally attributed to "breathing" of the telescope, i.e., the change in focus due to displacement of the secondary mirror in response to thermal cycling. If

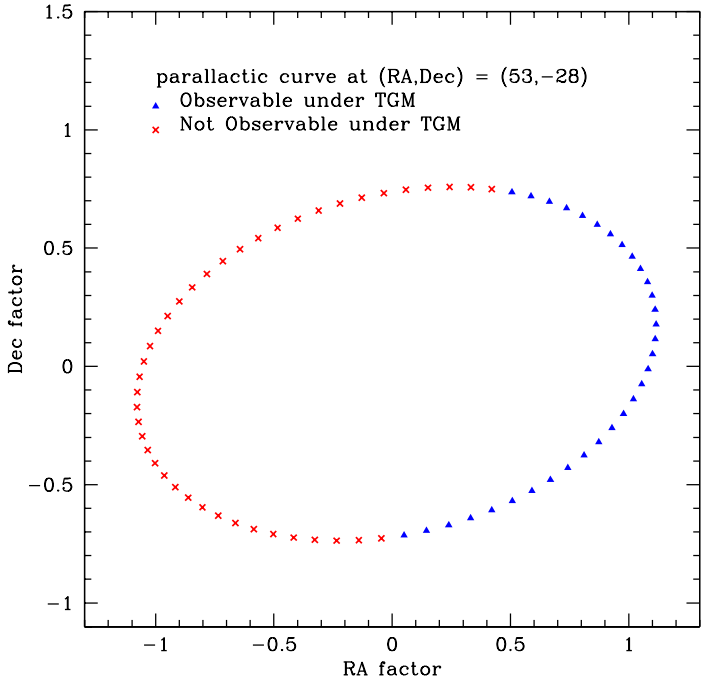

Figure 5: Schedulability of a target field as a function of the fields parallax factor, or day of year. Observation at epochs of "blue" parallax factors, or days, can be scheduled in two-gyro mode, the "red" days cannot. for clarity, every fifth day is plotted. The choice of the HST UDF is for illustration only.

not accounted for, this effect would impress an error of about 5 mas on FGS astrometric measurements. As discussed in the introduction, the effect is easily removed provided several selected stars are observed multiple times over the course the *HST* visit. The signature of this drift is, as expected, unchanged in two-gyro operations compared to three-gyro operations.

2.3. Transfer Mode

As discussed above, under two-gyro operations the spacecraft jitter and drift characteristics are not appreciably different than in three-gyro operations. Therefore, no change in the quality of data acquired by FGS1r in Transfer mode is expected. Indeed, the results of the February test demonstrated this. GO and calibration data obtained since the commencement of two-gyro operations in August 2005 confirm this conclusion.

3. Scheduling

Some astrometric programs may be impacted due to the more stringent scheduling constraints in two-gyro operations. In particular, the fixed head star trackers (FHTS) must not be occulted by the Earth since they are needed to coarse guide the telescope as the guiders acquire their respective guide stars (the FHSTs view the sky in a direction ≈perpendicular to the V1 axis). This eliminates access to one of the two epochs of maximum parallax factor for a given target field, an effect that has greater impact with decreasing declination. For illustration, Figure 5 shows the schedulability of observations in the HST UDF as a function of the field's parallax factors in RA and Declination, which each symbol corresponding to a particular day of year (every fifth day was plotted for readability). For this particular field,

the optimal times to observe for a parallax program would be at the extreme ends of the ellipse. However, in two-gyro mode observations can only be scheduled on the "blue" days.

The impact of these scheduling constraints will increase the error of a parallax measurement, but not generally by the amount which is proportional to the reduction in difference of the parallax factors at the epochs of the observations. Other sources of error enter as well, such as residuals of the geometric distortion calibration, and the conversion from relative to absolute parallax. Nonetheless, for targets near the distance limit (d $>\approx$ 400 pc) of accurate FGS parallaxes, the scheduling constraints become important for moderate to low declination fields.

4. Summary

FGS astrometry in two-gyro mode yields data of the same quality as that obtained in three-gyro mode. No modifications to the calibration pipeline or the calibration program are found to be necessary. The new scheduling restrictions can lead to a reduction in the accuracy of parallax measurements for distant objects at moderate to low declinations.

References

Nelan, E., et al. 2004, *Fine Guidance Sensor Instrument Handbook*, version 14.0 (Baltimore: STScI)

Nelan, E. 2005, *FGS Astrometry in the Feb 2005 On-orbit Two-Gyro Mode Test* www.stsci.edu/hst/HST_overview/TwoGyroMode/documents/FGS.pdf

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The Optical Field Angle Distortion Calibration of HST Fine Guidance Sensors 1R and 3

B. McArthur, G. F. Benedict and W. H. Jefferys

Astronomy Department, University of Texas, Austin, Texas 78712

E. Nelan

Space Telescope Science Institute, Baltimore, Maryland 21218

Abstract. To date five OFAD (Optical Field Angle Distortion) calibrations have been performed with a star field in M35, four on FGS3 and one on FGS1, all analyzed by the Astrometry Science Team. We have recently completed an improved FGS1R OFAD calibration. The ongoing Long Term Stability Tests have also been analyzed and incorporated into these calibrations, which are time-dependent due to on-orbit changes in the FGS. Descriptions of these tests and the results of our OFAD modeling are given. Because all OFAD calibrations use the same star field, we calibrate FGS 1 and FGS 3 simultaneously. This increases the precision of our input catalog,resulting in an improvement in both the FGS 1 and FGS 3 calibrations. A redetermination of the proper motions,using 12 years of HST data has significantly improved our calibration. Residuals to our OFAD modeling indicate that FGS 1 will provide astrometry superior to FGS 3 by $\sim 20\%$. Past and future FGS astrometric science supported by these calibrations is briefly reviewed.

1. Introduction

The largest source of error in reducing star positions from observations with the Hubble Space Telescope (HST) Fine Guidance Sensors (FGS's) is the Optical Field Angle Distortion (OFAD). Description of previous analyses can be found in McArthur et al. (2002, 1997), Jefferys et al. (1994), and Whipple et al. (1994,1996). The precise calibration of the distortion can only be determined with analysis of on-orbit observations. The Long Term STABility tests (LTSTAB), initiated in fall 1992, are an essential component of the OFAD calibration, and provide information on temporal changes within an FGS. They also provide indicators that a new OFAD calibration is necessary. This paper reports the results of the continuing OFAD calibration of FGS 3 and a newest OFAD calibration for FGS 1, including the LTSTAB tests. Past astrometry 2 and future astrometric results anticipated from FGS 1 are briefly reviewed.

2. Motivation and Observations

A nineteen orbit OFAD (Optical Field Angle Distortion) was performed in the spring of 1993 for the initial on-orbit calibration of the OFAD in FGS3. The first servicing mission made no changes to the internal optics of the three Fine Guidance Sensors (FGS) that are used for guiding and astrometry on HST. However, the subsequent movement of the secondary mirror of the telescope to the so-called "zero coma" position did change the morphology of the FGS transfer functions (Ftaclas et al. 1993). Therefore, a five orbit post servicing mission delta-OFAD calibration plan was designed and executed. After detection by the

LTSTAB of increasing incompatibility with the spring 1994 delta-OFAD calibration, an 11 orbit OFAD was performed in the fall of 1995 to recover the error budget for astrometry, after In the spring of 1997 a five orbit OFAD was performed on FGS3 after the second servicing mission. In December of 2000,a 14 orbit OFAD was performed on FGS1R, which replaced FGS3 as the prime astrometer for scientific observations. FGS1R, an enhanced FGS with an adjustable fold-flat mirror that can be commanded from the ground, had replaced the original FGS1 instrument in February of 1997 in SM2(Servicing Mission 2). Seventy LTSTABS (Long Term Stability Tests) have been performed in both FGS1R and FGS3 to assess time-dependent changes. A current list of the OFAD and LTSTAB tests is shown in Table 1.

3. Optical Field Angle Distortion Calibration and Long Term Stability Test

The Optical Telescope Assembly (OTA) of the HST (Hubble Space Telescope) is a Aplanatic Cassegrain telescope of Ritchey-Chrètien design. The aberration of the OTA, along with the optics of the FGS comprise the OFAD. The largest component of the design distortion, which consists of several arcseconds, is an effect that mimics a change in plate scale. The magnitude of non-linear, low frequency distortions is on the order of 0.5 seconds of arc over the FGS field of view. The OFAD is the most significant source of systematic error in position mode astrometry done with the FGS. We have adopted a pre-launch functional form originally developed by Perkin-Elmer (Dente, 1984). It can be described (and modeled to the level of one millisecond of arc) by the two dimensional fifth order polynomial:

$$\begin{aligned} x' = {} & a_{00} + a_{10}x + a_{01}y + a_{20}x^2 + a_{02}y^2 + a_{11}xy + a_{30}x(x^2+y^2) + a_{21}x(x^2-y^2) \\ & + a_{12}y(y^2-x^2) + a_{03}y(y^2+x^2) + a_{50}x(x^2+y^2)^2 + a_{41}y(y^2+x^2)^2 \\ & + a_{32}x(x^4-y^4) + a_{23}y(y^4-x^4) + a_{14}x(x^2-y^2)^2 + a_{05}y(y^2-x^2)^2 \end{aligned}$$

$$\begin{aligned} y' = {} & b_{00} + b_{10}x + b_{01}y + b_{20}x^2 + b_{02}y^2 + b_{11}xy + b_{30}x(x^2+y^2) + b_{21}x(x^2-y^2) \\ & + b_{12}y((y^2-x^2) + b_{03}y(y^2+x^2) + b_{50}x(x^2+y^2)^2 + b_{41}y(y^2+x^2)^2 \\ & + b_{32}x((x^4-y^4) + b_{23}y(y^4-x^4) + b_{14}x(x^2-y^2)^2 + b_{05}y(y^2-x^2)^2 \end{aligned} \tag{1}$$

where x, y are the observed position within the FGS field of view, x', y' are the corrected position, and the numerical values of the coefficients a_{ij} and b_{ij} are determined by calibration. Although ray-traces were used for the initial estimation of the OFAD, gravity release, outgassing of the graphite-epoxy structures, and post-launch adjustment of the HST secondary mirror required that the final determination of the OFAD coefficients a_{ij} and b_{ij} be made by an on-orbit calibration.

M35 was chosen as the calibration field. Since the ground-based positions of our target calibration stars were known only to 23 milliseconds of arc, the positions of the stars were estimated simultaneously with the distortion parameters. This was accomplished during a nineteen orbit calibration, executed on 10 January 1993 in FGS number 3. GaussFit (Jefferys, 1988), a least squares and robust estimation package, was used to simultaneously estimate the relative star positions, the pointing and roll of the telescope during each orbit (by quaternions), the magnification of the telescope, the OFAD polynomial coefficients, and these parameters that describe the star selector optics inside the FGS: ρ_A and ρ_B (the arm lengths of the star selectors A and B), and κ_A and κ_B (the offset angles of the star selectors). Because of the linear relationship between ρ_A, ρ_A , κ_A and κ_B, the value of κ_B is constrained to be zero. A complete description of that calibration, the analysis of the data, and the results are given in Jefferys et al. (1994).

Table 1: LTSTAB and OFAD Observations.

Orbit	Julian Date	Year	Day	FGS	Observation	Coefficient Set
1	2448959.340822	1992	337	3	LTSTAB	1
2	2448971.061435	1992	349	3	LTSTAB	1
3-21	2448997.782164	1993	10	3	OFAD	1
22	2449082.954086	1993	95	3	LTSTAB	1
23	2449095.742836	1993	108	3	LTSTAB	1
24	2449096.613044	1993	109	3	LTSTAB	1
25	2449226.341817	1993	238	3	LTSTAB	1
26	2449255.529236	1993	268	3	LTSTAB	1
27	2449283.771053	1993	296	3	LTSTAB	1
28	2449309.341898	1993	321	3	LTSTAB	1
29	2449379.838241	1994	27	3	LTSTAB	2
30	2449408.794850	1994	56	3	LTSTAB	2
31	2449437.560417	1994	85	3	LTSTAB	2
32	2449468.662153	1994	116	3	LTSTAB	2
33-37	2449469.602118	1994	117	3	Spring Delta-OFAD	2
38	2449593.554884	1994	241	3	LTSTAB	2
39	2449624.182975	1994	271	3	LTSTAB	2
40	2449652.274942	1994	299	3	LTSTAB	2
41	2449683.371435	1994	330	3	LTSTAB	2
42	2449711.665382	1994	359	3	LTSTAB	2
43	2449749.996910	1995	32	3	LTSTAB	2
44	2449780.160903	1995	62	3	LTSTAB	2
45	2449811.662894	1995	94	3	LTSTAB	2
46	2449838.070301	1995	120	3	LTSTAB	2
47	2449990.553542	1995	273	3	LTSTAB	3
48	2450018.625255	1995	301	3	LTSTAB	3
49	2450042.360197	1995	324	3	LTSTAB	3
50-60	2450052.674838	1995	335	3	Fall Delta-OFAD	3
61	2450112.122350	1996	29	3	LTSTAB	3
62	2450133.837824	1996	51	3	LTSTAB	3
63	2450158.835440	1996	76	3	LTSTAB	3
64	2450174.716192	1996	92	3	LTSTAB	3
65	2450199.778704	1996	117	3	LTSTAB	3
66	2450321.550822	1996	239	3	LTSTAB	3
67	2450353.777465	1996	271	3	LTSTAB	3
68	2450377.443275	1996	294	3	LTSTAB	3
69	2450416.366701	1996	333	3	LTSTAB	3
70	2450480.031933	1997	31	3	LTSTAB	3
71	2450518.768090	1997	70	3	LTSTAB	3
72-76	2450560.517523	1997	112	3	Spring Delta-OFAD	3
77	2450717.416169	1997	268	3	LTSTAB	3
78	2450743.225891	1997	294	3	LTSTAB	3
79	2450783.224190	1997	334	3	LTSTAB	3
80	2450822.077315	1998	8	3	LTSTAB	3
81	2450847.955266	1998	34	3	LTSTAB	3
82	2450904.886979	1998	91	1	LTSTAB	4
83	2450924.644942	1998	111	3	LTSTAB	3
84	2451054.361725	1998	240	3	LTSTAB	3
85	2451113.296366	1998	299	3	LTSTAB	3
86	2451121.224560	1998	307	1	LTSTAB	4
87	2451153.943299	1998	340	3	LTSTAB	3
88	2451163.019213	1998	349	1	LTSTAB	4
89	2451184.786771	1999	6	1	LTSTAB	4
90	2451189.556088	1999	11	3	LTSTAB	3
91	2451300.596829	1999	122	3	LTSTAB	3
92	2451300.664236	1999	121.2	1	LTSTAB	4
93	2451416.507917	1999	238	3	LTSTAB	3
94	2451430.269572	1999	251	1	LTSTAB	4
95	2451555.127963	2000	11	1	LTSTAB	4
96	2451555.199688	2000	11	3	LTSTAB	3
97	2451649.638229	2000	106	1	LTSTAB	4
98	2451653.660590	2000	110	3	LTSTAB	3
99	2451783.159410	2000	239	1	LTSTAB	4
100	2451830.321088	2000	286	1	LTSTAB	4
101-114	2451899.105289	2000	355	1	OFAD	4
115	2451968.923102	2001	59	1	LTSTAB	4
116	2452021.654896	2001	112	1	LTSTAB	4
117	2452137.970671	2001	228	1	LTSTAB	4
118	2452201.355764	2001	291	1	LTSTAB	4
119	2452263.961701	2001	354	1	LTSTAB	4
120	2452274.313264	2001	364	1	LTSTAB	4
121	2452295.219942	2002	20	3	LTSTAB	3
122	2452370.867882	2002	96	1	LTSTAB	4
123	2452384.694618	2002	110	1	LTSTAB	4
124	2452520.528970	2002	246	1	LTSTAB	4
125	2452581.074271	2002	306	1	LTSTAB	4
126	2452631.260370	2002	356	1	LTSTAB	4
127	2452700.778437	2003	0610	1	LTSTAB	4
128	2452749.608044	2003	110	1	LTSTAB	4
129	2452883.589838	2003	244	1	LTSTAB	4
130	2452953.9450698	2003	314	1	LTSTAB	4
131	2452997.011609	2003	357	1	LTSTAB	4
132	2453002.8817249	2003	363	1	LTSTAB	4
133	2453066.547211	2004	62	1	LTSTAB	4
134	2453118.7410301	2004	114	1	LTSTAB	4
135	2453234.513866	2004	234	1	LTSTAB	4
136	2453359.537685	2004	355	1	LTSTAB	4
137	2453363.535428	2004	359	1	LTSTAB	4
138	2453431.767731	2005	61	1	LTSTAB	4
139	2453480.676921	2005	110	1	LTSTAB	4
140	2453599.337049	2005	228	1	LTSTAB	4
141	2453639.847894	2005	269	1	LTSTAB	4
142	2453680.156250	2005	309	1	LTSTAB	4

In late fall 1992, just prior to the 1993 OFAD calibration, a series of one orbit long-term stability tests (LTSTAB) was initiated. These tests had two seasonal orientations, a spring orientation taken from an orbit of the OFAD, and a fall orientation, which was a 180 degree flip of the spring orientation. LTSTABs have been performed several times in each of the orientations, spring and fall, every year.

The LTSTAB is sensitive to scale and low order distortion changes. It is an indicator of the validity of the current OFAD coefficients and the need for recalibration. The LTSTAB series immediately showed that the scale measured by the FGS was changing with time. The indication of this change was seen in the large increase with time in the post-fit residuals from a solution that solved for a constant sets of star positions, star selector encoder (SSE) parameters, and OFAD parameters. The amount of scale change is too large to be due to true magnification changes in the HST optical telescope assembly. These changes could be due to water desorption in the graphite-epoxy components within the FGS. Initially the scale -like change was modeled by allowing a variation in the star-selector-A effective lever arm(ρ_A). Since 1995, the change has been modeled by allowing a change in both ρ_A and κ_A(the offset angle of the star selector).

A five orbit delta-OFAD was performed on 27 April 1994 after the first servicing mission to assess the distortion changes caused by the secondary mirror movement to the zero coma position. Significant effects in the OFAD (in addition to the scale-like changes) at the level of 10 mas were found. The LTSTAB tests have revealed continued permutations in the FGS. In addition to the scale changes, in mid-1995 we began to recognize higher order distortion changes. These changes manifested themselves as something that looks like a radial scale variation and is fairly well modeled by alterations in the third order terms in Eq. (1). We had also noted that the residuals from the fall orientation LTSTABS are consistently higher than for the spring in FGS3.

An eleven orbit delta-OFAD was performed in the late fall of 1995, to analyze temporal changes, and upgrade the y-axis coverage. The star catalog was redetermined with input from the three OFAD experiments of 1993, 1994 and 1995 to minimize the OFAD distortion that could have been absorbed by the catalog positions. A more complete analyses of this delta-OFAD can be found in McArthur 1997.

In the spring of 1997 a second servicing mission replaced FGS1. A five orbit delta-OFAD was performed in FGS3, repeating the orientation of spring 1994. The coefficients produced by this 5 orbit delta-OFAD did not provide a better calibration than the 11 orbit Fall of 1995 delta-OFAD calibration, so these orbits were used instead as LTSTABS. Two LTSTABS were performed in Spring 1997, one before and one after the second servicing mission. With scale and offset removed, a comparison yielded an rms of 0.965 mas, indicating stability of FGS3 across the servicing mission.

At the end of 2000, a 14 orbit OFAD was executed in FGS1R, for a total of approximately observations. Figure 1 shows the rotations and offsets of FGS1R in this OFAD calibration. In the first FGS1R OFAD, we entered the McNamara (1986) proper motion values as observations with error in a quasi-Bayesian fashion, instead of being applied as constants. More recently, we carefully examined the HST derived proper motions on their own, in cases where there were significant numbers of observations over significant time spans. These HST derived motions, are used in a weighted solution with the McNamara proper motions for the stars not so frequently observed. For this calibration, we ran a model which performed a simultaneous solution of OFAD polynomials, star selector encoder (sse) parameters, proper motions, drift parameters, and catalog positions. This model had over 15,000 equations of conditions using all 142 OFAD and LTSTAB plates. Only the OFAD plates determined the OFAD polynomials and complete sse parameters, while the LTSTAB combined with the OFAD plates contributed to a time-varying rhoA and kA, proper motions, and catalog positions. Each plate formed its own drift and rotation parameters. A systematic signature in the X residuals from the four OFAD analysis remains. This signature differs between FGS3 and FGS1. It appears as a very distinctive curve in the x

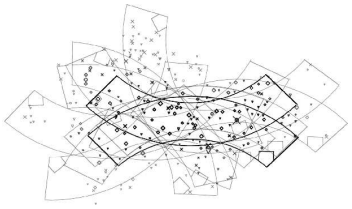

Figure 1: Rotation and Offsets of FGS1R Winter 2000 OFAD

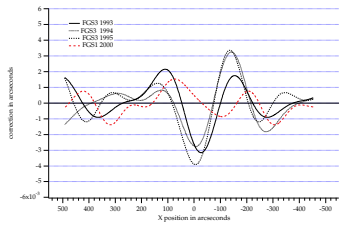

Figure 2: Four frequency Fourier series correction of systematic signature in X Residuals

component residuals as a function of position angle in the FGS field of view (Figure 2). The curve cannot be modeled by the fifth order polynomial. We have used a four frequency Fourier series to remove this effect. The size of this effect, in an RMS sense over the entire field of view of the FGS, is about one millisecond of arc. However, the peak-to-peak values near the center of the field of view can be as large as 7 mas in FGS3. The FGS1 systematic is much smaller with a peak to peak of about 2.5 mas. The source of this unexpected distortion is not yet known but it may be due to the way the FGS responds to the spherically aberrated HST beam.

On the basis of almost ten years of monitoring the distortions in FGS 3 we have concluded that at the level of a few milliseconds of arc, the optical field angle distortion in HST FGS 3 changes with time. These changes can be monitored and modeled by continuing the LTSTAB tests, which also alerts us to the need for a new OFAD calibration. There remains some dichotomy between the OFAD calibration data taken in the spring and that taken in the fall.

Five sets of OFAD coefficients (Eq. 1) and star selector parameters (M, ρ_A, ρ_B,κ_A and κ_B) have been derived for reductions of astrometry observations. The average plate residuals for these determinations are listed in Table 2. Comparisons of grids created with each set of FGS3 OFAD coefficients and distortion parameters indicate that the OFAD has changed around 10 milliseconds of arc in non-scalar distortion between calibrations (which have spanned 12-18 months)in FGS3.

Each LTSTAB is associated with a specific set of coefficients Table 1. In the boundary area betweeen two OFAD experiments, the observations are reduced with both sets of OFAD separately to determine which coeffecients produce the best ρ_A κ_A fit of the LTSTAB.

The values of ρ_A and κ_A determined by the LTSTABS and OFADS in FGSs 1 and 2 are illustrated in Figure 3,5, 4and 6. The error bars for these determination are smaller than the symbols. For reduction of science astrometry data, the ρ_A κ_A parameters are determined by interpolation of the two nearest LTSTABS in time.

The most recent analysis of the OFAD, which includes the HST determined proper motions, presents a significantly improved distortion calibration as seen in Table 2, listed as the New Winter 2000. The residual profile of the fgs1R stars is shown in Figure 7. The LTSTABS which have been performed since the implementation of two-gyro mode have been consistent with previous LTSTABS and we do not expect any deterioration in this mode of the calibration of the FGS.

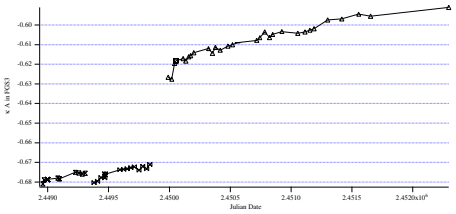

Figure 3: κ_A fit of the LTSTABS in FGS3

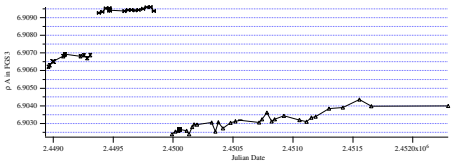

Figure 4: ρ_A fit of the LTSTABS in FGS3

Table 2: OFAD σ in millseconds of arc

OFAD	FGS	σ X	σ Y
Winter 2000	3	2.3	2.2
Winter 2000	1	1.6	1.7
New Winter 2000	3	1.5	1.8
New Winter 2000	1	1.3	1.3

In addition to the OFAD calibration, both Lateral Color and Cross Filter calibrations have been determined. Each FGS contains refractive elements. The position measured for a target star (relative to a differently colored set of reference stars) could depend on its intrinsic color, hence the lateral color corrections are applied. The large dynamic range in magnitudes (3 ? V ? 17) requires a neutral-density filter. There is a small shift in position (due to filter wedge) when comparing the positions of a bright star to a faint reference frame. This shift is determined in the Cross Filter calibration.

4. Past and Ongoing Astrometric Science with HST FGS

FGS 3 has been used to determine the first astrometrically determined mass of an extrasolar planet, which is around the star GL 876 (ApJ letters, in press). It has been used to obtain many trignometric parallaxes. Targets included distance scale calibrators (δCep - Benedict et al., 2002b; RR Lyr - Benedict et al., 2002a), interacting binaries (Feige 24 - Benedict et al., 2000), and cataclysmic variables (RW Tri - McArthur et al., 1999; TV Col- McArthur et al., 2001; SS Cyg, U Gem & SS Aur, Harrison et. al, 1999). It was also involved in an intensive effort to obtain masses and mass ratios for a number of very low-mass M stars (for example, GJ 22, GJ 791.2, GJ 623, and GJ 748- Benedict et al., 2001). The average parallax precision resulting from FGS 3 was $\sigma_\pi = 0.26$ mas.

FGS 1 has been used to determine the parallaxes of several cataclysmic variables (EX Hya, EF Eri, V1223 Sgr)(Beuermann et al., 2003, 2004; Harrison et al., 2004), parallaxes of a representative set of AM CVn stars, an independent parallax of the Pleiades (Soderblom et al., 2005), a parallax of the central star of NGC 6853 (Benedict et al., 2003) and the masses of extrasolar planets around 55 Cancri (McArthur et al., 2004) and ϵ Eridani. The masses of the extrasolar planets around υ Andromeda and additional planets are now being studied. A distance determination to the Cepheids is being calculated at present from HST

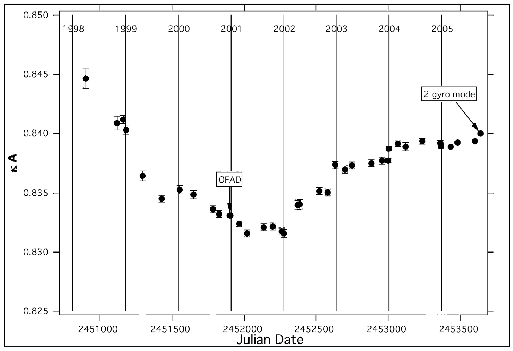

Figure 5: κ_A fit of the LTSTABS in FGS1

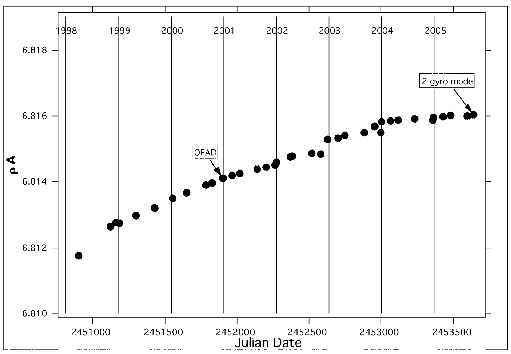

Figure 6: ρ_A fit of the LTSTABS in FGS1

parallaxes. FGS 1 is also involved in an ongoing effort to obtain masses and mass ratios for additional set of low-mass M stars.

A continued program of LTSTAB monitoring and OFAD updates is essential to the success of these long-term investigations with FGS 1.

5. Conclusions

We have shown that continued OFAD calibration of the Fine Guidance Sensors can reduce this source of systematic error in positions measured by the FGS's to the level of less than 2 mas. However, changes in the FGS units continue to occur, even twelve years after launch. These changes require periodic updates to the OFAD to maintain this critical calibration.

Acknowledgments. The Astrometry Science Team is supported by NASA NAG5-1603. We are grateful to Q. Wang for the initial modeling of the OFAD and D. Story and A.L. Whipple for their earlier contributions to this calibration. We thank L. Reed for her long-term contribution to our knowledge of FGS3. We thank Gary Welter and Keith Kalinowski for their interest and assistance at Goddard Space Flight Center. We thank all the members of the STAT, past and present for their support and useful discussions.

References

Benedict, G.F., Nelan, E., McArthur, B., Story, D., van Altena, W., Ting-Gao, Y., Hemenway, P.D., Shelus, P.J., Whipple, A.L., Franz, O.G., Fredrick, L.W., and R.L. Duncombe. 1993, PASP105, 487- 493

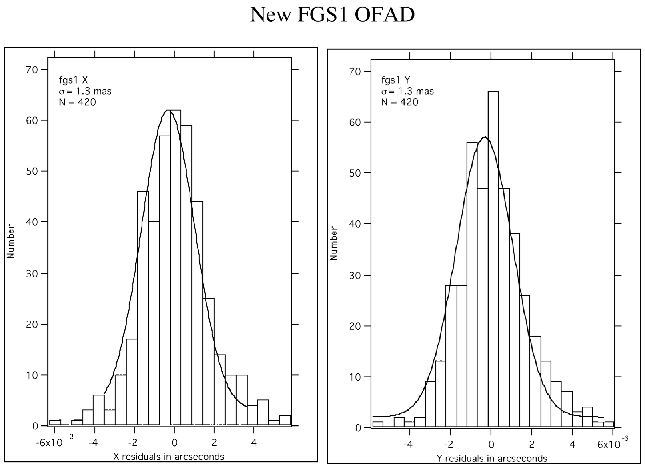

Figure 7: Residual Histrogram from OFAD calibration of FGS1R

Benedict, G. F., McArthur, B., Chappell, D. W., Nelan, E., Jefferys, W. H., van Altena, W., Lee, J., Cornell, D., Shelus, P. J., Hemenway, P. D., Franz, O. G., Wasserman, L. H., Duncombe, R. L., Story, D., Whipple, A., & Fredrick, L. W. 1999, AJ, 118, 1086

Benedict, G. F., McArthur, B. E., Franz, O. G., Wasserman, L. H., Nelan, E., Lee, J., Fredrick, L. W., Jefferys, W. H., van Altena, W., Robinson, E. L., Spiesman, W. J., Shelus, P. J., Hemenway, P. D., Duncombe, R. L., Story, D., Whipple, A. L., & Bradley, A. 2000, AJ, 119, 2382

Benedict, G. F., McArthur, B. E., Franz, O. G., Wasserman, L. H., Henry, T. J., Takato, T., Strateva, I. V., Crawford, J. L., Ianna, P. A., McCarthy, D. W., Nelan, E., Jefferys, W. H., van Altena, W., Shelus, P. J., Hemenway, P. D., Duncombe, R. L., Story, D., Whipple, A. L., Bradley, A. J., Fredrick, L. W. 2001, AJ, 121, 1607

Benedict, G. Fritz, McArthur, B. E., Fredr ick, L. W., Harrison, T. E., Lee, J., Slesnick, C. L., Rhee, J., Patterson, R. J ., Nelan, E., Jefferys, W. H., van Altena, W., Shelus, P. J., Franz, O. G., Wass erman, L. H., Hemenway, P. D., Duncombe, R. L., Story, D., Whipple, A. L., & Br adley, A. J. 2002A, AJ123, 473

Benedict, G. F., McArthur, B. E., Fredrick, L. W., Harrison, T. E., Slesnick, C. L., Rhee, J., Patterson, R. J., Skrutskie, M. F., Franz, O. G., & Wasserman, L. H. 2002B, AJ124, 1695

Benedict, G. F., McArthur, B. E., Fredrick, L. W., Harrison, T. E., Skrutskie, M. F., Slesnick, C. L., Rhee, J., Patterson, R. J., Nelan, E., Jefferys, W. H., van Altena, W., Montemayor, T., Shelus, P. J., Franz, O. G., Wasserman, L. H., Hemenway, P. D., Duncombe, R. L., Story, D., Whipple, A. L., & Bradley, A. J. 2003, AJ, 126, 2549

Beuermann, K., Harrison, T. E., McArthur, B. E., Benedict, G. F., & Gänsicke, B. T. 2004, A&A, 419, 291

Beuermann, K., Harrison, T. E., McArthur, B. E., Benedict, G. F., & Gänsicke, B. T. 2003, A&A, 412, 821

Bradley, A., Abramowicz-Reed, L., Story, D., Benedict, G., and W. Jefferys. 1991, PASP, 103, 317

Dente, G. and Abramowitz-Reed, L. Space Telescope Optical Field Angle Distortion Calibration Algorithm. 1984, Perkin-Elmer ST-SE-3336.

Ftaclas, C., Nonnenmacher, A., Weindling, F., Story, D., and E. Nelan. 1993, Ap. Opt. 32, 1696

Harrison, T. E., McNamara, B., Szkody, P., McArthur, B. E., Benedict, G. F.,Klemola, A., & Gilliland, R. L. 1999, ApJ, 515, L93

Harrison, T. E., Johnson, J. J., McArthur, B. E., Benedict, G. F., Szkody, P., Howell, S. B., & Gelino, D. M. 2004, AJ, 127, 460

Jefferys, W., Fitzpatrick, J., and McArthur, B. 1988 Celest. Mech. 41, 39.

Jefferys, W.J., Whipple, A., Wang, Q., McArthur, B., Benedict, G.F., Nelan, E., Story, D, and L. Abramowicz-Reed. 1994, in Calibrating Hubble Space Telescope, J.C. Blades and S.J. Osmer (Baltimore: STScI) 353-374.

McArthur, B. E., Benedict, G. F., Jefferys, W. H., and Nelan, E. 1997, in The 1997 *HST Calibration Workshop : Hubble after the Installation of the ACS and the NICMOS Cooling System*, Santiago Arribas, Anton Koekemoer, and Brad Whitmore (Baltimore: STScI) 373.

McArthur, B. E., Benedict, G. F., Jefferys, W. H., and Nelan, E. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 472.

McArthur, B. E., Benedict, G. F., Lee, J., Lu, C. -L., van Altena, W. F., Deliyannis, C. P., Girard, T., Fredrick, L. W., Nelan, E., Duncombe, R. L., Hemenway, P. D., Jefferys, W. H., Shelus, P. J., Franz, O. G. & Wasserman, L. H. 1999, ApJ, 520, L59

McArthur, B. E. Benedict, G. F., Lee, J., van Altena, W. F., Slesnick, C. L., Rhee, J., Patterson, R. J., Fredrick, L. W., Harrison, T. E., Spiesman, W. J., Nelan, E., Duncombe, R. L., Hemenway, P. D., Jefferys, W. H., Shelus, P. J., Franz, O. G., Wasserman, L. H., 2001, ApJ, 560, 907

McArthur, B. E., Endl, M., Cochran, W. D., Benedict, G. F., Fischer, D. A.;,Marcy, G. W., Butler, R. P., Naef, D., Mayor, M., Queloz, D.,Udry, S., Harrison, T. E., 2004 ApJ614, L81

McNamara, B., & Sekiguchi, K., 1986. AJ91, 557

Soderblom, D.R.,Nelan, E., Benedict, G. F., McArthur, B., Ramirez, I., Spiesman, W., Jones, Burton F., 2005, AJ129, 1616S

van Altena, W. F., Lu, C. -L., Lee, J. T., Girard, T. M., Guo, X., Deliyannis, C., Platais, I., Kozhurina-Platais, V., McArthur, B., Benedict, G. F., Duncombe, R. L., Hemenway, P. D., Jefferys, W. H., King, J. R., Nelan, E., Shelus, P. S., Story, D., Whipple, A., Franz, O. G., Wasserman, L., Fredrick, L. W., Hanson, R. B., Klemola, A. R., Jones, B. F., Mendez, R., Tsay, W. -S., Bradley, A. 1997. ApJ Lett, 486, 123

Wang, Q., and Jefferys, W. 1990, in Satatistical Analysis of the Optical Field Angle Distortion of the Hubble Space Telescope Using the Bootstrap Method, Austin:Univesrity of Texas, Department of Astronomy Institute, 1-99.

Whipple, A., Jefferys, W., Wang, Q., McArthur, B., Benedict, G.F., Nelan, E., Story, D., and L. Abramowicz-Reed. 1994, in Calibrating Hubble Space Telescope, J C. Blades and S.J. Osmer ed(s)., (Baltimore: STScI) 375-379.

Whipple, A., McArthur, B., Wang, Q., Jefferys, W., Benedict, G.F.,Lalich, A.M., Hemenway, P.D., Nelan, E., Shelus, P.J., and Story, D.. 1995, in Calibrating Hubble Space Telescope:Post Servicing Mission, A. Koratkar and C. Leitherer ed(s)., (Baltimore: STScI) 119-127.

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

HST Temporal Optical Behavior: Models and Measurements with ACS

R. B. Makidon, S. Casertano, and M. Lallo

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. While HST provides a stable stellar image relative to ground-based observatories, it features its own characteristic changes in the PSF it delivers to the Science Instruments (Lallo et al. 2005). HST focus has been monitored and adjusted throughout the life of the observatory. More recently, the resolution and off-axis location of ACS/HRC has allowed us to accurately measure changes in coma and astigmatism as well. The aim of this current work is to relate the accurate and reliable phase retrieval measurements of wavefront error back to characterizations more common to science data analysis (e.g. encircled energy, FWHM, ellipticity). We encourage the further examination of these effects on HST science observations.

1. Introduction

It has been known since early in the Mission that the focal length of HST varies on both orbital and longer time scales (Bély 1993). These changes have generally been attributed to a physical motion of the secondary mirror (SM) resulting from variations in the metering truss structure that supports it. HST focus has always been monitored on at least a monthly basis (Figure 1). The primary purpose for this monitoring is to characterize the focus state of the observatory and accurately plan the time and amount of the occasional SM move for focus maintenance.

ACS/HRC has enhanced our ability to measure the additional image aberrations of coma and astigmatism. Like focus, these aberrations vary over the HST orbit, suggesting the possibility of a more complex motion of the SM, or other sources of misalignment. As part of our routine monitoring program in recent years, phase retrieval analysis (Krist & Burrows 1995) is performed on stellar targets in HRC and values of focus, 0 and 45 degree astigmatism, and x and y coma (Zernike coefficients 4 through 8) are obtained over an orbit at roughly monthly intervals.

The plots in Figure 2 show these data. Units are in microns rms wavefront error for coma and astigmatism, and microns of SM despace

2. PSF Morphology in ACS

It is well known that the morphology of the ACS PSF in both WFC and HRC exhibits appreciable field dependence (see Krist 2003). However, the understanding of how that PSF morphology varies over time and as a function of e.g. focus has not been well understood. Understanding the ACS PSF and the factors that affect it has become increasingly important as science observations in fields such as weak lensing and circumstellar environments continue to push the limits of what is observable.

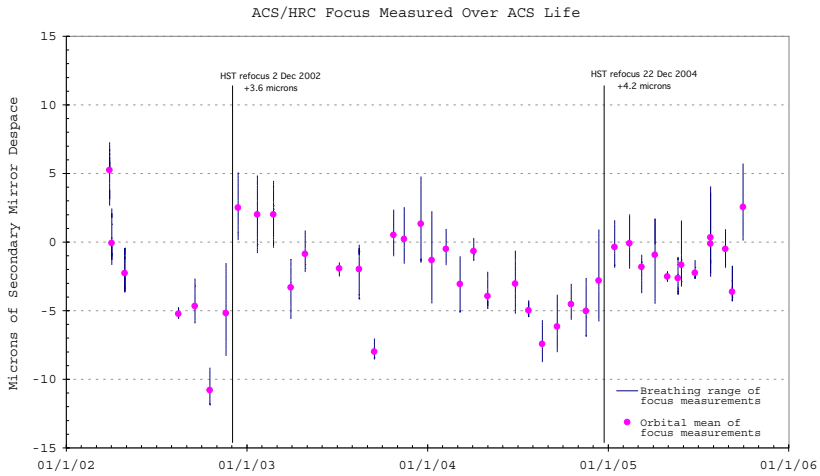

Figure 1: Secondary mirror (SM) despace (in microns) determined from a monthly series of observations of isolated bright stars in the ACS/HRC. Phase Retrieval was used to first calculate the Zernike focus term (z4) for the star in each observation, which can then be related to microns of motion at the SM. The filled circle represents the orbital mean of the focus observations, while the bar represents the range of measurements obtained for that orbit.

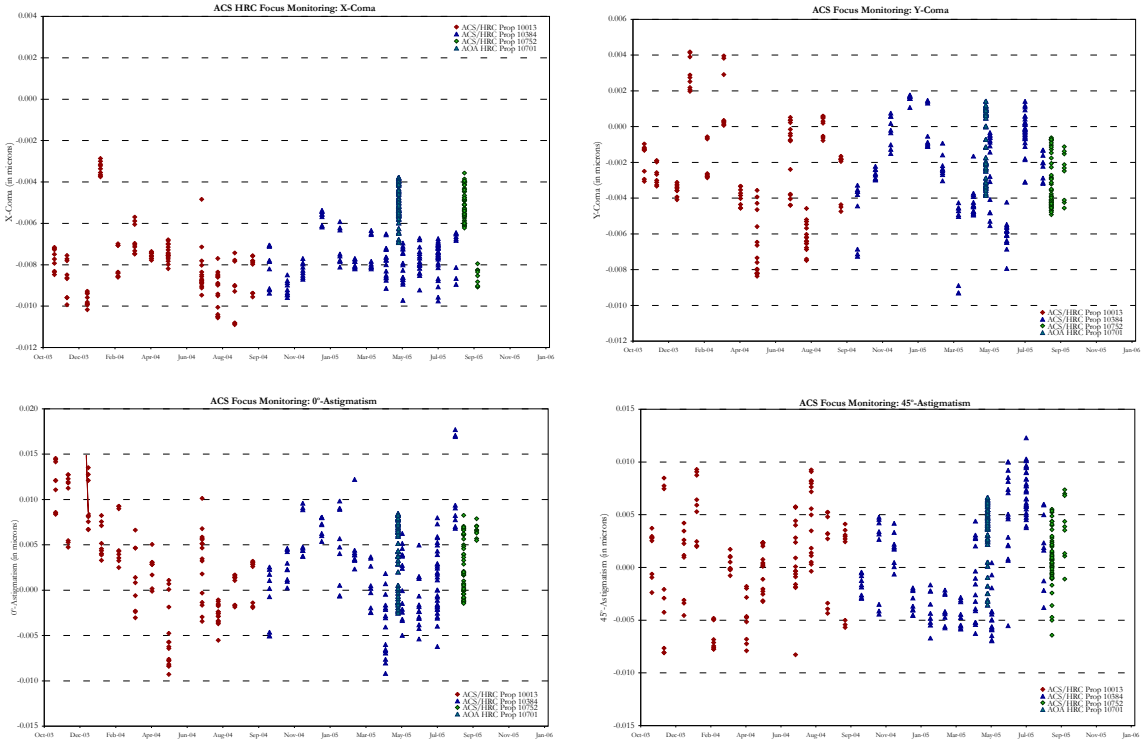

Figure 2: Coma (top row) and astigmatism (bottom row) measured from the monthly set of observations used to monitor ACS and HST focus. Phase Retrieval was used to determine the Zernike terms z5 through z8 from these data. Each point corresponds to an observation within an orbit, and are plotted in units of microns *rms* of wavefront error as a function of time.

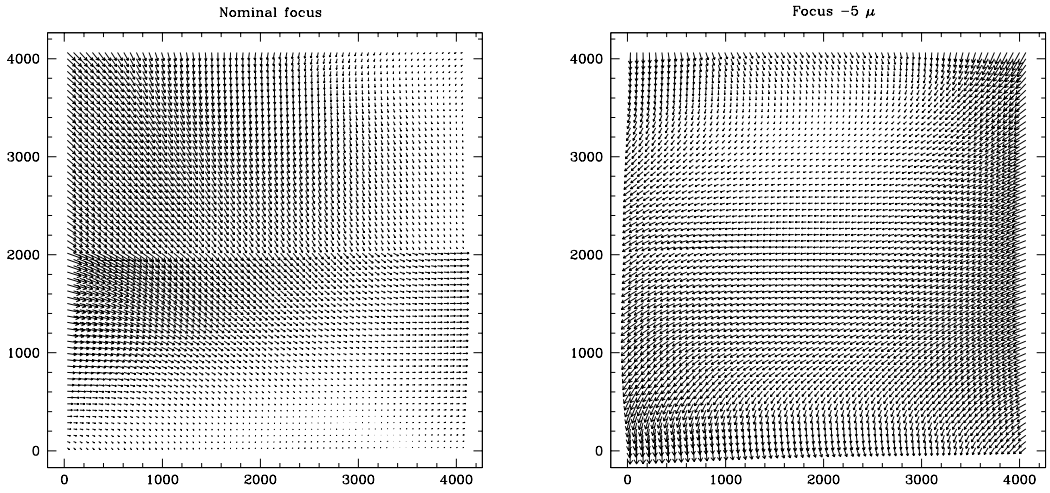

Figure 3: PSF ellipticities in the ACS/WFC measured from Tiny Tim PSF models. At left: ellipticities measured from models at nominal focus. At right: ellipticities measured from models at $\sim -5\mu$m of SM despace.

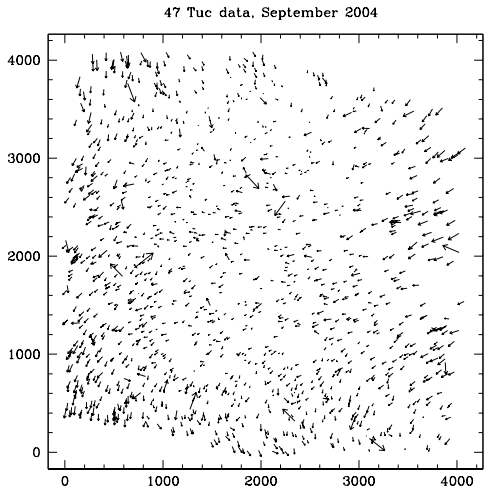

Figure 4: PSF ellipticities of stellar sources in 47 *Tuc* observed in the ACS/WFC with *F*814*W* while *HST* was known to be in a negative focus state of $\sim -5\mu$m

2.1. Observed and Modeled PSF Ellipticities in WFC

Using Tiny Tim, we generated model PSFs on a finely-sampled grid of field positions on the WFC (Figure 3). Here, measured PSF ellipticities across the WFC field are shown for nominal focus (left) and a Secondary Mirror despace of -5μm or 30nm wavefront error (middle). In Figure 4 we show PSF ellipticity measurements using stellar data from observations of 47 *Tuc* observed with *F*814*W* while *HST* was known to be in a negative focus state of $\sim -5\mu$m (Sept 2004). The ellipticities measured in the 47 *Tuc* data are in good agreement with those in the -5μm model. It is clear that even such a modest despace (typical of orbital focus variation) results in an observable change in ellipticity over the WFC field.

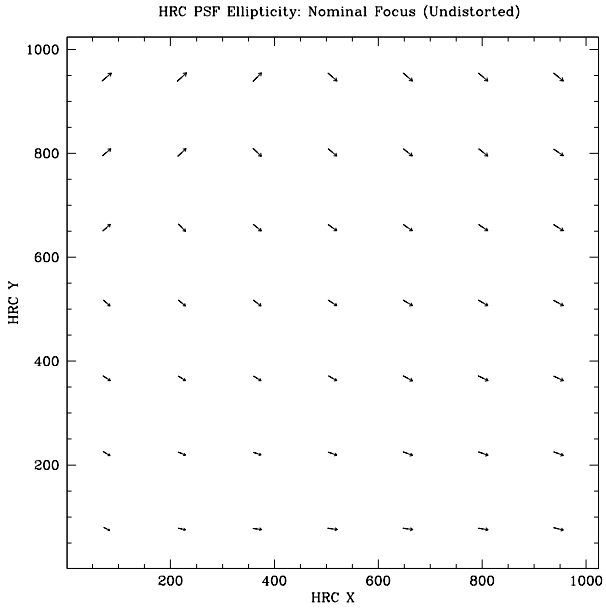

Figure 5: PSF ellipticities in the ACS/HRC measured from Tiny Tim models at nominal focus. The length of the vector represents $10e$, where e is the ellipticity of the PSF.

2.2. Modeled Ellipticities in HRC

Using Tiny Tim, we generated model PSFs with the F555W filter on a 7×7 grid of positions across the HRC field. We examined the variation of the PSF ellipticity as a function of focus (holding coma and astigmatism constant), and as a function of coma and of astigmatism at nominal focus. We show the measured ellipticities (calculated from image moments) for the nominal focus case (Fig.5), and the difference in ellipticities at 5μm on either side of nominal focus (Fig.6).

Fig.7 illustrates the change in ellipticities over typical ranges of coma (left pane) and astigmatism (right) at nominal focus. It is clear that as with WFC, changes in focus induce a measurable effect on ellipticity, presumably due to astigmatic changes in the PSF due to the combined effects of the HST OTA and the ACS optics. However, neither coma nor astigmatism variations alone at a constant focus (as might be due to slight motions of optics within the HRC instrument itself) have shown a comparably significant effect on the observed PSF morphology.

3. Conclusion

Changes in the field-dependent morphology of the PSF in both HRC and WFC appears more strongly sensitive to focus variations due to HST SM despace than to the observed range of coma and astigmatism. While these aberrations have been shown to exhibit behavior on orbital and longer timescales, the causes of these variations are not yet well-understood. Though the use of phase retrieval can precisely quantify these variations, more familiar—though less sensitive—"real world" characterizations of the observed PSF (e.g. encircled energy, FWHM, and ellipticity) are far less sensitive to the range of coma and astigmatism observed than to focus. Understanding the state of the HST focus at the time of one's observations is necessary to achieve the photometric and astrometric precisions required of many current HST programs (see Gilliland et al. 2002; Suchkov & Casertano 1997), and to accurately describe the morphologies of barely resolved objects.

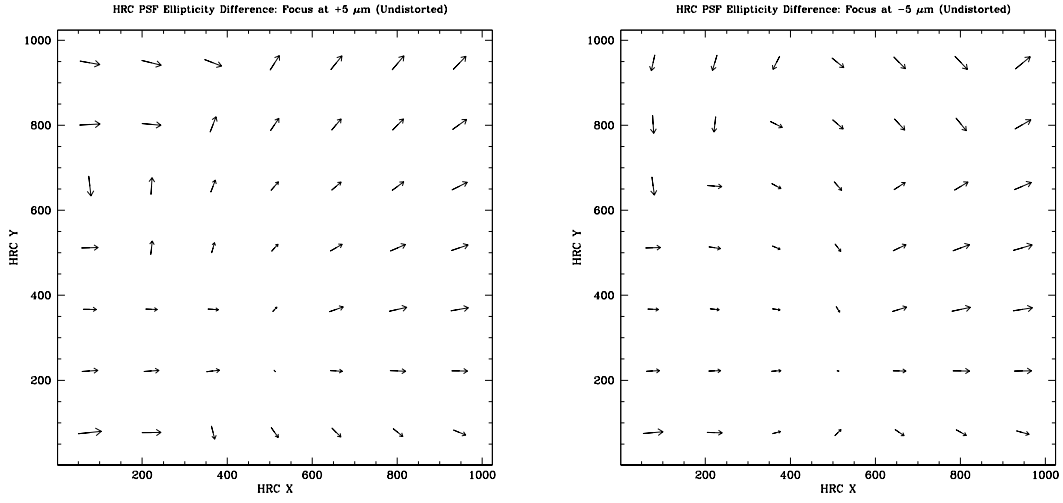

Figure 6: Measured changes in PSF ellipticity in the ACS/HRC from Tiny Tim models at focus positions $\pm 5\mu$m on either side of nominal focus. Changes are shown relative to ellipticities measured from Tiny Tim models at nominal focus, with the length of each vector given by $10(\Delta e)$, where Δe is the difference in measured ellipticity relative to nominal focus at that field position.

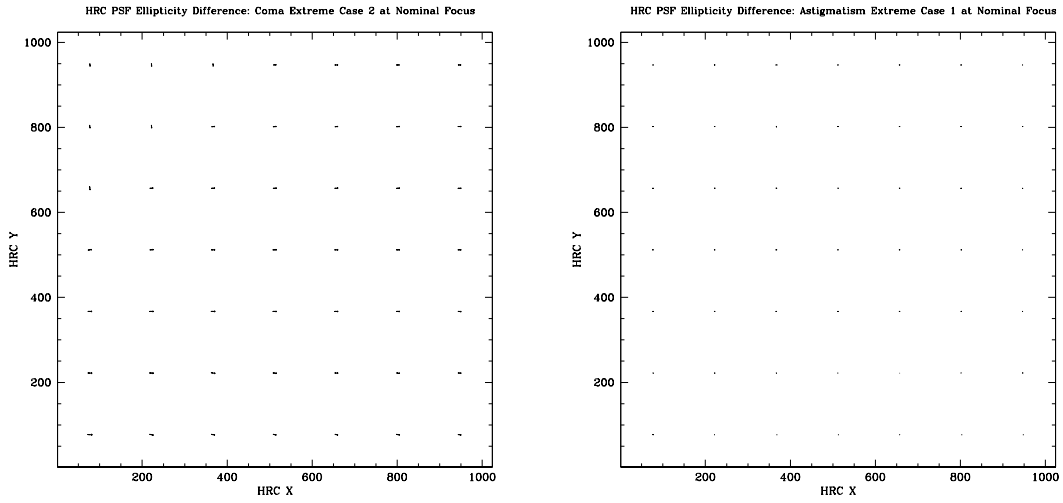

Figure 7: Measured changes in PSF ellipticity in the ACS/HRC from Tiny Tim models at nominal focus over typical ranges of coma (left) and astigmatism (right). Changes are shown relative to ellipticities measured from Tiny Tim models at nominal focus for default coma and astigmatism values. Here, the length of each vector is given by $10(\Delta e)$, where Δe is the difference in measured ellipticity relative to default coma and astigmatism values at zero focus.

References

Bély, P., Hasan, H., & Miebach, M., 1993, "Orbital Focus Variations in the Hubble Space Telescope", SESD-93-16

Gilliland, R.L. et al., 2000, ApJ, 545, L47

Krist, J. E. & Burrows, C. J., 1995, Appl. Opt. 34, 495

Krist, J. E., 2003, "ACS WFC & HRC field-dependent PSF variations due to optical and charge diffusion effects", *Instrument Science Report* ACS 2003-06 (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Lallo, M., Makidon, R. B., Casertano, S., Gilliland, R., & Stys, J., 2005, "HST Temporal Optical Behavior & Current Focus Status", *Instrument Science Report* TEL 2005-03 (Baltimore: STScI), available through http://www.stsci.edu/hst/

Suchkov, A. & Casertano, S., 1997, "Impact of Focus Drift on Aperture Photometry", *Instrument Science Report* WFPC2 97-01 (Baltimore: STScI), available through http://www.stsci.edu/hst/wfpc2

Part 6. GSC-II and Data Analysis

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

The New GSC-II and its Use for HST

B. McLean

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. The second generation HST Guide Star Catalog (GSC-II) provides a significant improvement over the original GSC-I. It provides ICRS-based positions (International Celestial Reference System) with improved accuracy for almost a billion objects as faint as J=21. This will become the default reference frame for HST observations starting in cycle 15 and will reduce the absolute errors of the astrometry in the observation headers from 1-2" down to 0.3-0.5". In addition, the new astrometry is being incorporated into the Digitized Sky Survey (DSS) images and can also be used to improve the astrometric accuracy of archival HST data.

1. Introduction

The current Guide Star Catalog (GSC-I) used for HST pointing is becoming less and less accurate as the effects of proper motion continue to increase the relative position errors between the stars. These motions have gradually increased the guide star (GS) acquisition failure rate as the positions 'age' since the epoch of the original measurement. In addition, HST proposals required that the target coordinates were in the same GS reference frame which increased the complexity of preparing phase 2 submissions. Recognizing that these issues (and others) would eventually have to be addressed, STScI in collaboration with other institutions (McLean et al 2004), embarked on a project to produce a 2nd generation GSC that would improve the coordinates and provide additional observation planning capabilities.

2. GSC-I Limitations

The GSC-I was constructed in the 1980's using photographic Sky Survey plates from the Palomar and UK Schmidt telescopes. A complete description of the project is published in Lasker at al 1990 (Paper 1), Russell et al 1990 (Paper 2) and Jenkner et al 1990 (Paper 3). At that time, there was was no single homogeneous all-sky astrometric reference catalog with the required accuracy so three different reference catalogs were used (the AGK3 for plates 0°to +90°, SAOC -60°to 0°, CPC -90°to -60°). Since there are significant discontinuities between the catalogs this requires GOs to ensure that supplied target coordinates are in the reference frame defined by the GSC-I. In order to simplify the task of observation planning and measuring coordinates, STScI arranged the distribution of the GSC-I and the all-sky Digitized Sky Survey (DSS) images to the community on CDROM and later on-line access via the world wide web.

The sky survey plates were taken between 1975 and 1987 and the GSC-I objects have measured coordinates between these epochs. The expected increase in relative position errors due to proper motions were factored into the GS error budget based on the HST pointing accuracy and the size of the (original) science instrument (SI) apertures. Over the last few years, we have noted an increase in the number of failed observations for small apertures due to proper motions in the GS and the targets. This was temporarily alleviated

when STIS failed (since there was no longer an SI with a small aperture), however if COS is installed in a future servicing mission then we would be unable to meet the requirements for placing targets in its small apertures.

The GSC-I was originally designed to have GS in the range 9-15 over the entire sky that could be used by the FGS to point HST. This was to efficiently select GS in a highly automated way. Once STIS (and later ACS) was selected as a replacement instrument, it was recognized that it would be necessary to check the FOV for each pointing to ensure that no UV-bright star ($V \leq 20$m) would be in the aperture of the MAMA detectors. A 'bright' star would result in detector damage. This became a 'Health and Safety' issue requiring a visual inspection of the field using the DSS for each pointing. It was quickly realized that a deeper object catalog would allow similar efficiency improvements to this "Bright Object Protection' task.

3. GSC-II Overview

Even though the concept of a second-generation catalog was developed soon after the original GSC-I was completed it took many years to get all the necessary components in place. The most fundamental item was to scan more recent epoch plates. This was accomplished by obtaining access to the POSS-II surveys in the northern hemisphere and in collaboration with the AAO, arranging for a new southern hemisphere survey. The STScI scanning machines were significantly upgraded to complete digitization of all the plates at the required resolution (Laidler et al 1996). In order to proceed with the construction, digitization and distribution of the new DSS and the development and production of the GSC-II (McLean et al 1998), STScI partnered with a large number of institutions to cost-share and distribute the work in return for early access to the datasets for telescope operations.

One of the early design decisions was to scan and process all the available Palomar and UKSTU surveys in order to give both multi-passband and multi-epoch observations to allow the measurement of proper motions and colors of the objects wherever possible. This provides at least 3 passbands and 2 epochs (including a recent epoch) everywhere on the sky. The images were processed to the plate limits to get the maximum amount of information for observation planning. All of these observations are stored in a 3TB object-oriented database (Greene 1998) which is mined to produce an export catalog with the derived parameters for almost a billion objects.

The availability of homogeneous all-sky astrometric reference catalogs (ACT, TYCHO-2) based on the ICRS reference frame was a major factor in the astrometric improvements along with improved plate modeling and reduction techniques. The result is that GSC-II has an absolute astrometric error of 0.25" (1-sigma) over the entire sky. In order to obtain a photometric calibration fainter than 15th magnitude a major observing campaign was begun to obtain CCD frames for every survey field to create a network of secondary reference stars for calibrating the plates. Despite using photographic material we typically have a photometric error of 0.3mag (1-sigma) and a classification rate (star/non-star) of 95%.

4. HST Observing

Whilst the fundamental goal of this change is to improve the precision of HST pointing, the GO will see some simplification of the procedures used to provide target coordinates.

4.1. Phase 2 Changes

A new keyword (ICRS) has been added to the phase 2 proposal preparation options for the coordinate reference frame. The existing keywords are still present for cycle 15 whilst

this transition takes place, but ICRS is the default selection and recommended whenever possible. Many of the recent major surveys and catalogs are ICRS-based (2MASS, SDSS, USNO-B, FIRST etc), so coordinates may be taken directly from these sources without transferring them into the GSC reference frame.

4.2. Measuring DSS Images

If your target is an extended non-stellar source such that it's catalogued coordinate in GSC-II (or any other catalogue) is not where you wish to place the SI aperture, and yet is bright enough to appear on the DSS images you can measure the coordinates as before. The GSC-II coordinates are primarily based on the IIIaF (red) surveys. This is the most recent epoch so we strongly recommend using these images to minimize the errors. The DSS images headers have also been updated to include the ICRS astrometry. Whilst the old GSC-I calibration keywords are still in the headers (for backward compatibility), a new set of GSC-II calibration keywords and FITS extensions are also included. This calibration however, is a 3-step non-linear transformation that most FITS readers do not yet implement. To provide the ICRS information we have modified the code running on the DSS server at STScI to dynamically compute a simple FITS WCS (World Coordinate System) calibration that all packages should be able to utilize. This uses the local GSC-II solution at the position of each image extraction to compute the WCS keywords which it then inserts into the returned FITS image. You can measure the positions using any analysis tool that utilizes the FITS WCS. Once your measured position is entered into your proposal, please select the ICRS reference frame option.

4.3. Measuring CCD images

If your target is not visible directly on the DSS images or otherwise catalogued, you are expected to provide coordinates by transferring the GSC reference frame to your images (this includes HST images). Instead of using GSC-I, one can now use the much deeper GSC-II, *or any other ICRS-based catalogue* which makes finding reference stars easier.

4.4. Tools

Access to the GSC-II and DSS images will be built into the cycle 15 release of the Astronomers Proposal Tool (APT) which provides a convenient interface using the Visual Target Tuner (VTT). Alternatively, one will be able to find links to web forms available from the STScI website. If a GO has successfully observed a target with HST in a previous cycle and is confident that the target has zero proper motion, we also provide a web form that can convert your GSC-I based coordinate to the GSC-II ICRS coordinate. This tool measures the *average* coordinate shift for the GSC-I objects over the HST field of view and corrects the supplied coordinate accordingly.

5. Status

At the time of submitting this manuscript, HST has successfully used the GSC-II for a first on-orbit test. Assuming that the remaining tests are successful, then the ICRS/GSC-II will become the *default* reference system for HST observing in cycle 15 and beyond.

We have successfully completed a feasibility study (Koekemoer et al 2005) of using the GSC-II to recalibrate HST ACS images, and looking ahead, STScI is considering a project to do this for as many images in the HST archive as possible. This would significantly improve the scientific usefulness of HST images when comparing them to observations at different wavelengths. In the case of images without sufficient GSC-II stars to recalibrate, then updating the coordinates using the GSC-I to GSC-II offset will still improve the astrometry and put it on the standard ICRS.

Even further in the future, it is planned to use the GSC-II (with improvements and cross-matching to 2MASS for IR magnitudes) as the basis for JWST operations. Studies have established that it is deep enough to provide sufficient GS (even at near-IR wavelengths) to point JWST.

Acknowledgments. We acknowledge the contributions of our many collaborating institutions. These include: Osservatorio Astronomico di Torino, Space Telescope - European Coordinating Facility, European Space Agency - Astrophysics Division, European Southern Observatory, International GEMINI Project, Canadian Astronomical Data Center, Centre de Donne Stellaire, Anglo-Australian Observatory, Royal Observatory Edinburgh, National Astronomical Observatory Japan, Beijing Astronomical Observatory, Palomar Observatory and the Leicester Database and Archive Service.

References

Greene,G., McLean,B., Lasker,B., Wolfe,D., Morbidelli,R. & Volpicelli,A. 1998, IAU Symposium 179, 474

Jenkner,H., Lasker,B.M., Sturch,C.R., McLean,B.J., Shara,M.M. & Russell,J.L. 1990, AJ, 99, 2081

Koekemoer,A.M., McLean,B., McMaster,M. & Jenkner,H. 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 417

Lasker, B.M., Sturch,C.R., McLean,B.J., Russell,J.L., Jenkner,H. & Shara,M.M. 1990, AJ, 99, 2019

Laidler,V., Sturch,C., Greene,G., Lasker,B.M., McLean,B., White,R., Bucciarelli,B., Lattanzi,M., Spagna,A. & Zacchei,A.2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI)1996, BAAS, 188, 5421L

McLean,B., Hawkins,G., Spagna,A., Lattanzi,M., Lasker,B., Jenkner,H. & White,R. 1998 IAU Symposium 179, 43

McLean,B., Lattanzi, M., Greene, G., Loomis, C. & Meakes, M. 2004, SpaceOps 2004 conference proceedings

Russell,J.L., Lasker,B.M., McLean,B.J., Sturch,C.R. & Jenkner,H. 1990, AJ, 99 2059

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Improving the Absolute Astrometry of HST Data with GSC-II

A. M. Koekemoer, B. McLean, M. McMaster and H. Jenkner

Space Telescope Science Institute, 3700 San Martin Dr., Baltimore MD 21218, USA

Abstract. We have used the new GSC-II to significantly improve the absolute astrometric accuracy of a large number of ACS datasets from the first years of ACS operations, reducing the initial levels of uncertainty (typically 1 - 2", or more in some cases) to an accuracy of about 0.1 - 0.3". There are typically enough GSC-II objects in the field of view of the ACS/WFC detectors to enable a robust solution to be determined by identifying the GSC-II objects and comparing their measured positions with those in the GSC-II. The resulting set of improved header coordinates are accurate enough to permit immediate multi-band comparisons between data from HST and other space-based and ground-based facilities, as well as a range of wavebands including X-ray, mid-IR, and radio data. The Institute has also begun planning to extend this work by including it in the automatic processing pipeline for ACS and eventually perhaps also for other HST instruments.

1. Introduction

A crucial scientific application of science data obtained with HST is to provide cross-comparison with other observations of the same targets, obtained with HST or different telescopes (ground-based or space-based). There are many examples where major scientific discoveries have only been made possible by combining multi-wavelength observations of the same targets. This is important not only in terms of the physics of the sources since many of them emit across a broad range of wavelengths (for example, active galactic nuclei which are detected in the radio, optical, and X-ray regimes), but also in studies of distant sources at very high redshift, where the emission from the sources can shift all the way from the optical/UV into the mid-infrared. In addition to multi-waveband comparisons, accurate absolute astrometry is also important to enable HST observations of adjacent or overlapping fields to be combined into a single contiguous image, as well as allowing observations of the same target from different dates to be combined and achieve greater depth, and to search for supernovae and other variable sources.

However, a fundamental limitation inherent in any given HST dataset is the fact that the absolute astrometric information in the headers is typically inaccurate by $\sim 1-2''$, thus $\sim 10-20$ times worse than the exquisite resolution of HST. As a result, significant amounts of work often need to be invested by researchers wishing to compare a given HST dataset to any other dataset, in order to ensure scientifically meaningful and accurate results.

This study is aimed at using sources identified in HST images to improve the absolute astrometry by about an order of magnitude, to $\sim 0.1-0.3''$, thus approaching the resolution of the telescope, and better than typical ground-based or other space-based facilities. This provides a significant improvement in the scientific quality of the data.

2. Sources of Astrometric Inaccuracies in HST Data

A major source of uncertainty in the absolute astrometry of HST data arises from the Guide Star Catalog I (GSC-I; Lasker et al. 1990; Russel et al. 1990; Jenkner et al. 1990) which has been used in the HST operational system for all cycles up to Cycle 14. The astrometry for any given observation is derived by assuming the location of the guide star to be accurately represented by its coordinates in the GSC-I. Subsequently, the telescope calculates an offset between the coordinates of the guide star and the target, and performs a slew that would place the target exactly at the correct location on the instrument aperture if the guide star position is accurate; the astrometry of the image is subsequently calculated on this basis. If the guide star is at a slightly different location on the sky, the astrometric header for the image remains unchanged, but the target is not placed at the correct position on the detector and the resulting astrometry for the entire image is incorrect.

The typical level of uncertainty of guide star positions in GSC-I is $\sim 1''$ although some objects can deviate by up to $\sim 2-3''$ (or more). These uncertainties arise from a combination of factors, including uncertainties in the plate solutions describing the distortion of the photographic plates, nearby stars, plate artifacts or other confusing sources that can affect the measured position, as well as the lack of proper motion information for guide stars in GSC-I, which can place them at a different position than expected.

An additional source of uncertainty arises from the location of the instrument apertures and the Fine Guidance Sensors (FGS) relative to one another in the HST focal plane. It has been demonstrated that over long periods of time, the instrument and FGS apertures can drift as a result of long-term changes in the structure of the telescope. These shifts can be comparable to the arcsecond-level uncertainties in the GSC-I; although they can be corrected on a routine basis with aperture updates, they nevertheless represent a significant additional source of uncertainty in determining the absolute astrometry of an image, even if the guide star position is known with complete accuracy.

3. General Methodology to Improve Absolute Astrometric Accuracy

The first step toward improving the astrometric accuracy begins with using the revised version of the Guide Star Catalog II (GSC-II; McLean et al. 2004; McLean 2006), which has greatly improved astrometry ($\sim 0.25''$ r.m.s.) and also contains far more objects, at least $\sim 5 \times 10^8$. In addition, GSC-II contains well-measured proper motions for all stars for which this is an important effect. Thus, the absolute astrometry of any given image can be immediately improved to some extent by examining the change in coordinates of its primary guide star from GSC-I to GSC-II. Starting in Cycle 15, GSC-II will be used in HST operations and will thus automatically produce more accurate astrometry for new data, but for older data this change needs to be applied retro-actively.

However, updating the guide star positions does not address any potential uncertainty in the instrument aperture locations, and in addition the final astrometry in the image would still be dependent upon the position of a single guide star. Although the r.m.s. positional accuracy of GSC-II is better than GSC-I, individual stars can still deviate by several times the r.m.s, thus up to $1''$ or more.

Therefore, a more direct means of improving the absolute astrometry of an image is to identify all the GSC-II sources that are on the image, and use their collective positions to reduce the contribution of errors from any single source to the overall astrometric correction. Not only does this provide a more accurate position if all the objects are well-behaved, but it provides additional protection against being adversely affected by a single bad object, since such objects can be rejected iteratively.

A crucial consideration in identifying GSC-II objects on the images is that the image should cover a sufficient area of sky to ensure the presence of a sufficient number of such

sources. If an exposure is too shallow or if the field of view is too small to contain any such objects, then the only available options are to update the astrometry based on the improved guide star position, along with the best knowledge of the location of the instrument in the telescope focal plane. This can be tracked with time and is the subject of another study.

4. The Present Study

In this study we consider primarily the technique of using GSC-II objects identified on the images themselves to improve the astrometry. Given the space density of GSC-II objects across the sky ($\sim 12,000$ per square degree, on average, to the sensitivity limits of the plates), the best instrument on HST to provide a demonstration of this technique is the Advanced Camera for Surveys / Wide Field Channel (ACS/WFC), which covers $\sim 3'$ on a side and therefore will contain ~ 30 GSC-II sources on average. In practice, fields at high galactic latitudes generally contain $\sim 5-10$ sources on a given ACS/WFC image, while fields near the galactic plane may contain up to 100 or more sources in extreme cases.

After demonstrating this technique for ACS/WFC, it can then be applied to data from any of the other instruments that have a sufficient number of GSC-II sources in the field (i.e., STIS, WFPC2 and NICMOS, as well as ACS/HRC and ACS/SBC), although such datasets will likely only be a fraction of the full range of datasets for these instruments since they have much smaller fields of view than ACS/WFC.

In order to provide a good demonstration of the general applicability of this technique, it needs to be shown to work for a wide range of different types of targets, through different filters, and using different observing strategies. We chose to examine the first year of ACS/WFC data (all of which are now public), in all the broad-band and medium-band filters, for all exposure times and observation strategies. The initial search for all ACS/WFC datasets in the HST archive yielded a total of 20,010 datasets (with an average of 3 exposures in each dataset). For the present study we excluded data from SM3B, calibration, engineering and parallel exposures (although the latter may be included in a future study), and also excluded several large programs (including GOODS, UDF, and COSMOS), for which the science teams themselves have already produced improved absolute astrometry based on extensive ground-based imaging. The resulting sample of all broad-band and medium-band prime ACS/WFC external science exposures yielded a total of 2,492 datasets obtained during the first year of ACS operations.

5. Technique Used to Improve Absolute Astrometry

There are fundamentally two different types of exposures in our final dataset:

1. A single exposure of a given field

2. Exposures that form part of a multi-exposure dataset

As a first step, we chose to concentrate on the second category of datasets, namely those with more than one exposure on a given field. The fundamental reason for this is that cosmic rays play a significant role in producing spurious source detections on an image, since several thousand of them are typically present on any exposure longer than a few minutes. Therefore, we found that astrometric registration of single-exposure images, while also possible in principle, required significantly more fine-tuning than in the case of multi-exposure datasets where the exposures could be combined to reject cosmic rays and produce a clean catalog. Single-exposure datasets will be discussed in a future study, while for the present demonstration we discuss astrometric improvements made to multi-exposure datasets.

Multi-exposure datasets were defined as those satisfying the following criteria, in order to allow satisfactory combination:

1. executed within the same visit, on the same target
2. using the same camera (ACS/WFC)
3. using the same filter

This definition allowed exposures to be grouped together if they were part of a pre-specified dither pattern (i.e., reproducing the current rules for creating associations from dither patterns), but in addition allowed exposures to be grouped together if they had been specified using POS TARGs or no dithering at all, yet satisfied the given criteria.

Each multi-exposure dataset was then run through MultiDrizzle (Koekemoer et al. 2002), which automatically aligned the individual exposures, made a clean "median" image, and then used the clean image to perform cosmic ray rejection before using Drizzle (Fruchter & Hook 2002) to combine all the input images onto a single, clean output image.

A catalog was then created from the clean output image. We investigated both SExtractor (Bertin & Arnouts 1996) as well as the IRAF DAOFIND software (Stetson 1987), and chose the latter since we were primarily interested in detecting stars, for which DAOFIND is able to provide more accurate centroiding. The DAOFIND parameters were then set to produce a catalog of stars on each drizzled, combined image, yielding their R.A. and Dec. positions as computed from the image header astrometry keywords. This catalog was then cross-correlated with the relevant section of the GSC-II, allowing an initial tolerance of 5″ and refining the positional offset iteratively until the shifts converged. In each case, the free parameters solved for included shifts in R.A. and Dec. We also initially considered solving for the small rotational offsets that can be introduced as a result of the guide star uncertainties. However, there are typically not enough GSC-II sources on a given image to allow rotations to be reliably determined, therefore we focus on solving for the shifts in R.A. and Dec., which represent the dominant uncertainty in the absolute astrometry.

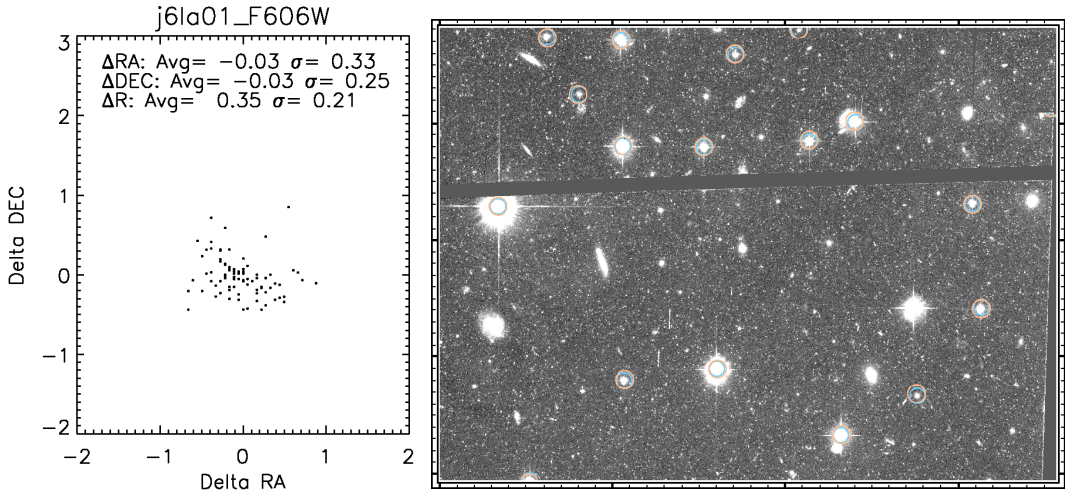

Figure 1: (Left) Example of the typical degree of scatter (in arcsec) in the comparison of measured vs. catalog positions for GSC-II objects on a representative drizzled image, after having removed the net offset in R.A. and Dec. originally present in the header astrometry, which for this image was $\sim 1.7''$. (Right) A 60″-wide portion of the same image, showing the identified GSC-II objects after the net shift in R.A. and Dec. had been removed. Circles mark the positions of sources in the catalog compared with their measured locations on the image.

Once the shift had been solved for the drizzled image, the values were then propagated back to the distorted frame of each input exposure (the "FLT" file), since ultimately the goal of this project is to improve the absolute astrometry of each exposure. However, for the purposes of this study, we present the offsets in R.A., Dec. for each drizzled image, since this is most directly relevant to the purpose of this study.

6. Results

After running DAOfind and automatically updating the drizzled image header coordinates with the resulting shift in R.A. and Dec., the new measured coordinates for all GSC-II objects on the image were computed and compared with their coordinates in the GSC-II catalog. In Figure 1 we show a typical plot, for an image that contained 86 matched sources from the GSC-II. This shows the typical spread in the locations of a given source, as well as the fact that some significant outliers generally exist even for GSC-II, although the peak of the distribution is very well behaved, with a 1-σ scatter of only $\sim 0.2 - 0.3''$. The outliers are often fainter objects near the plate limit, since we are using the full GSC-II catalog in carrying out these comparisons. However, it is encouraging to note that many of the non-stellar GSC-II sources are also very well behaved, even though they are often $\sim 2 - 3$ magnitudes fainter than the typical guide star magnitudes. All the drizzled images and plots were visually inspected to ensure that the algorithm had computed the correct shifts in all cases, and that image artifacts and other problems were not causing mis-registrations.

In Figure 2 we present the measured shifts for the full set of 152 combined, drizzled broad-band images in our sample, which altogether contain a total of 834 exposures. Each combined dataset represents an association (either resulting from a dither pattern or a set of POS TARG exposures with common filters), and contains exposures from the same visit, hence with the same guide stars. Therefore, the R.A. and Dec. shifts obtained from the drizzled image can be applied to all the input exposures comprising that image, since the same guide star was used on all the exposures with at most a re-acquisition (e.g., exposures obtained during different orbits of the same visit). It can be seen that less than 35% of the exposures had an original astrometric error less than $0.5''$, while more than about half of the exposures had an original astrometric error $> 1''$, and $\sim 25\%$ had an error $> 2''$. After applying the astrometric corrections, we re-measured the locations of all sources in the images compared with their GSC-II coordinates. This showed that the majority of the images have have an astrometric error $< 0.25 - 0.5''$, and only 10% have an error $> 0.75''$. Hence, these corrected images approach the intrinsic positional accuracy of GSC-II.

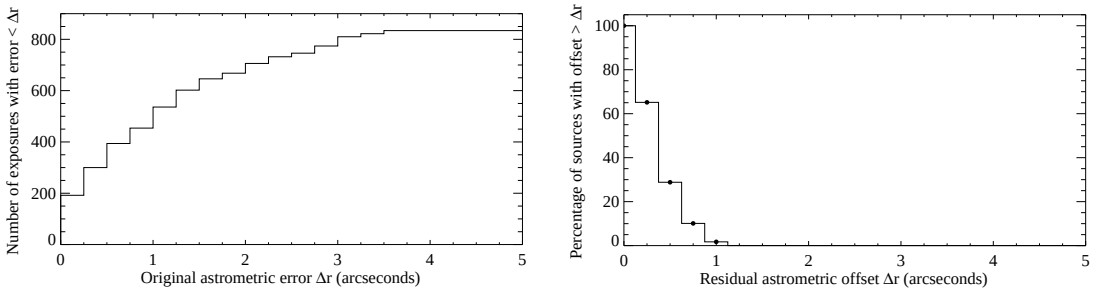

Figure 2: (Left) Cumulative distribution of the number of exposures that have an original astrometric error less than Δr, as a function of the original astrometric error Δr, thus representing the initial astrometric error of each image. (Right) Histogram showing the measured positions of sources in the images after correcting the image WCS, relative to their GSC-II coordinates. The majority of the sources now have an astrometric error $< 0.25{-}0.5''$, hence these corrected images approach the intrinsic positional accuracy of GSC-II.

7. Summary

We have presented a demonstration of a technique to improve the absolute astrometric accuracy of HST images, by producing a catalog for the image and comparing the resulting objects with those in the GSC-II on the same portion of the sky. This initial demonstration study was carried out on associated broad-band prime ACS/WFC imaging science observations obtained during the first year of ACS operation, consisting of a total of 152 drizzled combined datasets (representing a total of 834 separate exposures), covering a wide range of different types of science images. By utilizing the full GSC-II, we are generally able to obtain sufficient objects on each image to enable good shifts to be computed, typically improving the absolute astrometry to levels of $\sim 0.1-0.3''$, thus yielding an order of magnitude improvement over the original uncertainties inherent in the data.

Future follow-up work will include extending the images to include single exposures (which have more cosmic rays in them, and thus require more fine-tuning of the matching algorithms), as well as examining the feasibility of this technique on other HST cameras (ACS/HRC, ACS/SBC, WFPC2, NICMOS, STIS) which have smaller fields of view but which may contain GSC-II objects in some fraction of the images. Finally, by allowing all HST data to be tied to GSC-II, this study sets the stage for the eventual transition to GSC-II in operations, by allowing archival data to possess the same level of astrometric accuracy as future HST observations. By allowing much more accurate registration to data from other epochs, telescopes and wavebands, this represents a significant improvement in the scientific quality and discovery potential of HST data.

Acknowledgments. We are grateful to Rodger Doxsey and Ron Gilliland for useful discussions during the course of this work. The Guide Star Catalog was produced at the Space Telescope Science Institute under U.S. Government grant. These data are based on photographic data obtained using the Oschin Schmidt Telescope on Palomar Mountain and the UK Schmidt Telescope.

References

Bertin, E. & Arnouts, S. 1996, "SExtractor: Software for source extraction", A&AS 117, 393

Fruchter, A. S. & Hook, R. N. 2002, PASP, 114, 144

Jenkner, H., Lasker, B. M., Sturch, C. R., McLean, B. J., Shara, M. M. & Russell, J. L. 1990, AJ, 99, 2081

Koekemoer, A. M., Fruchter, A. S., Hook, R. N., & Hack, W. 2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 337

Lasker, B. M., Sturch, C. R., McLean, B. J., Russell, J. L., Jenkner, H. & Shara, M. M. 1990, AJ, 99, 2019

McLean, B., Lattanzi, M., Greene, G., Loomis, C. & Meakes, M. 2004, SpaceOps 2004 conference proceedings

McLean, B. 2006, The 2005 *HST Calibration Workshop*. Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 413

Russell, J. L., Lasker, B. M., McLean, B. J., Sturch, C. R. & Jenkner, H. 1990, AJ, 99 2059

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

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Multidrizzle and Tweakshifts: Overview and Future Plans

A. M. Koekemoer[1], A. S. Fruchter[1], R. N. Hook[2], W. Hack[1], C. Hanley[1]

Abstract. The MultiDrizzle software provides a flexible, unified interface to the various steps involved in registering, cleaning and combining dithered *HST* data. MultiDrizzle can be run on ACS, WFPC2, NICMOS and STIS imaging data in a single step, with default parameters chosen to provide good results for the majority of datasets if standard recommended dithering patterns were followed. Alternatively, the various parameters can be adjusted and specific steps can be run one at a time for additional flexibility, if desired for certain scientific applications. MultiDrizzle has also been incorporated into the *HST* OTFR pipeline for ACS data, thereby allowing automatic delivery of cleaned, combined data that can be used directly for science in many cases, or alternatively used as a baseline for re-running MultiDrizzle off-line using customized parameters, if required for certain scientific applications. Here we review the basic fundamentals of dithering, together with the current status of MultiDrizzle and describe plans for future related pipeline enhancements, including the Tweakshifts shift-refinement script.

1. Introduction

A widespread technique in obtaining Hubble Space Telescope *(HST)* observations involves the use of "dithering", or spatially offsetting the telescope in order to move targets to a number of different locations on the detector. There are a number of different scientific drivers for employing dithering, with different benefits/trade-offs depending on the size and layout of the dither steps and the scientific goals of the observations, as described in detail in the *HST Dither Handbook* (Koekemoer et al. 2002a). Here we review these considerations, and describe current progress with the MultiDrizzle software (Koekemoer et al. 2002b) which has been developed to provide a flexible, unified interface to the various steps involved in registering, cleaning and combining dithered *HST* data.

One of the most common reasons for dithering is to mitigate the relatively severe undersampling of the *HST* point-spread function (PSF) by the pixels in most of the detectors onboard *HST*. The undersampling imposes strong limitations on the accuracy with which morphological and photometric measurements can be obtained from exposures at any given single location. By obtaining a series of exposures shifted by non-integer pixel offsets, the information can be sampled on scales that are finer than the detector pixel size, thereby permitting the construction of a more finely sampled image when the individual exposures are combined. This is the motivation behind the "drizzle" software (Fruchter & Hook 2002), which is called by MultiDrizzle and makes use of linear reconstruction algorithms to allow input images to be resampled onto a finer output pixel grid.

In addition, the sensitivity of each pixel varies as a function of position across the pixel, hence the undersampling can lead to problems when pursuing very accurate stellar photometry since the measured flux of a star may vary significantly depending on how it is

[1]Space Telescope Science Institute, 3700 San Martin Dr., Baltimore, MD 21218, USA

[2]Space Telescope – European Coordinating Facility, Karl-Schwarzschild-Str. 2, D-85748 Garching, Germany

centered on a pixel. Sub-pixel dithering helps mitigate this intra-pixel sensitivity variation by allowing the pixel response function to be sampled more completely, thereby resulting in more robust photometric measurements.

Furthermore, dithering on larger scales to move targets to different locations on the detector can help improve the signal-to-noise ratio (S/N) by averaging out variations in flat field sensitivity, including variations in sensitivity from one pixel to the next as well as large-scale background variations, that may lie below the level of accuracy of the flat fields and other reference files. If no dithering is performed, then these uncertainties can act as systematic errors in any determination of photometric properties, hence they can be mitigated by averaging them out when the targets are moved around to different locations on the detector.

Another reason for dithering is to move bad pixels or detector blemishes around to different parts of the sky, so that any given object will not be too severely affected by them. This also includes moving the gaps that may be present between multiple detectors in the same instrument, for example the two chips of the Advanced Camera for Surveys / Wide Field Channel (ACS/WFC), or the four separate cameras of the Wide Field Planetary Camera 2 (WFPC2). This is often done when it is desirable to obtain reasonably uniform coverage across the field. Although the resulting depth in the gaps may be slightly shallower than the surrounding coverage, this is often preferable to having no coverage at all.

Finally, very large-scale offsets on the scale of the detector itself can be used if it is necessary to cover a contiguous region that is significantly larger than the detector size, thereby creating a "mosaic" of the target field. Strictly speaking, such offsets address a different class of issues than the smaller-scale dither offsets discussed above, but they are included in the discussion since many of the same technical considerations apply to them as well.

2. Dither Strategies in Practice

When planning to combine dithered data with MultiDrizzle, it is important to first plan the observations in a way that will maximize the scientific return. An important issue that should be taken into account when considering dither patterns is the degree of geometric distortion across the detectors, which is particularly severe for ACS/WFC and ACS/HRC that have about a $\sim 7\%$ plate scale change across the detector, in addition to substantial skew. The plate scale change means that a relatively small offset at the center, for example 7 pixels, gradually changes to a different offset across the chip until it contains an additional non-integer component of ~ 0.5 pixels at the corner. Similarly, a larger offset of 14 pixels at the center would gradually change to an additional 0.5 pixels halfway to the corner, then continue changing back to integer sampling by the time the corner is reached. Even larger dithers introduce additional cycling between sub-pixel and integer offsets between the center and corner of the chip. In reality it is somewhat more complicated since the ACS distortion contains a significant skew term in addition to the scale change, which means that the degree of non-integer pixel sampling varies across the detector in an even more intricate pattern for any given dither offset.

This has implications when considering the design of a dither pattern to choose for a given science program. If retaining absolutely uniform sub-pixel sampling across the entire detector is the most important consideration, for example in very critical photometric studies, and if some loss of coverage can be tolerated due to bad pixels and chip gaps, then it is plausible that dither offsets may be chosen that are entirely less than 1 pixel in extent. In this case the largest differential change in sampling introduced by the $\sim 7\%$ geometric distortion across the detector would correspond to ~ 3.5 milliarcseconds for the ACS/WFC and ~ 1.7 milliarcseconds for the ACS/HRC, which becomes comparable to the pointing uncertainty of *HST*. Such a strategy would not move bad columns or chip gaps around to

sufficiently different locations, hence resulting in a loss of coverage at those locations, but this may be tolerated if uniform sub-pixel sampling is critically important.

A more common strategy involves combining sub-pixel offsets with larger integer pixel offsets, to ameliorate the effects of hot pixels and bad columns as well as the ACS/WFC chip gap. In this case, the $\sim 2.5''$ offset along the y-axis that is needed to cover the chip gap corresponds to ~ 50 pixels, and thus will most certainly introduce variable amounts of sub-pixel sampling along the y-direction, particularly if multiple exposures are used that cover the gap in increments of 50 pixels. This is generally mitigated by using a dither pattern that fully samples 4 different pieces of sub-pixel phase space, such as the pre-defined 4-point BOX dither, or a 2-point primary LINE dither with an additional 2-point secondary LINE dither at each of the two pointings. By obtaining a sufficient number of dither pointings which cycle through sub-pixel and integer pixel sampling in different ways across the detector, it is generally possible to ensure reasonably uniform sub-pixel sampling across the entire field. This strategy is commonly used for extragalactic field studies where the targets are generally resolved, where it is important to obtain some degree of sub-pixel sampling, but equally important to move hot pixels and bad columns to different locations, as well as to cover the ACS/WFC chip gap. We note that these issues are not as critical for WFPC2 which only has $\sim 1-2\%$ distortion; however, one of its cameras (the PC) has about half the pixel scale of the other three cameras, thereby necessitating a special dither pattern optimized to provide sub-pixel sampling in all the cameras simultaneously.

Generally, the 4-point dither patterns which include half-pixel offsets provide sufficient sub-pixel sampling for almost all the observing programs on *HST*. This is because the degree of undersampling by the pixels is typically no more than about a factor of 2 at worst for ACS/WFC, while the ACS/HRC provides sampling that is essentially Nyquist-limited toward the red part of the spectrum. Furthermore, the ACS/WFC requires a few minutes to read out one exposure, which places a natural limit on the number of separate dither positions that can be obtained in one orbit for small to medium-sized programs. In addition, the typical pointing uncertainty of *HST* is in the range $\sim 2-3$ milliarcseconds (Gilliland et al. 2005; Koekemoer et al. 2005a, and these proceedings), or $\sim 25\%$ of a half-pixel dither with ACS/HRC, thus it becomes infeasible to expect dither patterns to execute perfectly if they attempt to sample scales much finer than 1/2 or 1/3 of an HRC pixel. Some of the deeper observing programs, such as the Hubble Ultra Deep Field, used 12-point patterns (effectively a 4-point box pattern replicated 3 times, with slight additional shifts, randomized to some extent by the *HST* pointing). The primary effect of this is to ensure more uniform sub-pixel sampling across the entire field, particularly when adding in the small random offsets of a few milliarcseconds that are a consequence of the guide star re-acquisitions over large numbers of orbits.

3. MultiDrizzle - Combining Dithered Images

Once a series of dithered images have been obtained, several steps need to be carried out in order to combine them. First, their relative shifts need to be determined, for which the *HST* header astrometry is generally accurate enough to be used, if the exposures were all obtained in the same visit using the same guide stars. The typical offset and pointing accuracy of *HST* is $\sim 2-3$ milliarcseconds when there are no guide star problems, which is thus at the level of 0.1 pixel for ACS/HRC and 0.05 pixel for ACS/WFC (and 0.02 pixel for the WFPC2/WF chips), and generally below the level of other effects such as PSF-related changes. Occasional guide star problems can increase these errors, and certain science programs demand more accurate alignment, as well as cases where data need to be combined that were taken in different visits using different guide stars (in which case their astrometry could be offset by up to a few arcseconds). For such applications we describe

later the Tweakshifts script, which can directly use the contents of the images themselves to improve their alignment.

After the image shifts have been determined, they need to be geometrically transformed to a set of common, registered output images where all the sources are coincident in pixel space. This permits the creation of a clean approximation to the final image, by evaluating quantities such as the median value of all pixels at a given location, which mitigates the contribution from outliers such as cosmic rays or hot pixels. Once this clean image has been created, it is then transformed back to the frame of each of the input exposures. Comparing the clean image with each input exposure then allows the identification of cosmic rays and bad pixels, which can be incorporated into a pixel flag mask. The pixel masks are then used in the final "drizzle" step, which maps all the input exposures onto the output frame, and combines them by performing a weighted sum of their pixel values, using the pixel masks to eliminate the contribution from rejected pixels.

The MultiDrizzle script was developed to provide a flexible, unified interface to all the above steps. Prior to the availability of MultiDrizzle, all these steps needed to be carried out manually using a variety of different IRAF scripts, with a significant amount of detailed book-keeping required to ensure that all the correct intermediate files were present. This was burdensome even for small and intermediate programs, and became prohibitive for large programs. The goal of MultiDrizzle is to remove the overhead of keeping track of all the intermediate information, while still retaining the ability to run each of these steps separately. It provides enough parameters to permit detailed control of the behavior of each step if required, while also providing default values for these parameters that allow it to be run in a "one-touch" mode, which can start with a set of calibrated exposures and automatically produce a cosmic-ray cleaned, distortion corrected, drizzled combined output image, by invoking a single command.

Data that were obtained using standard recommended dither patterns can be provided to MultiDrizzle to carry out all the steps of sky subtraction, registration, cosmic ray rejection and final drizzle combination all in a single command, simply by specifying the list of input files, for example:

```
--> multidrizzle input=*_flt.fits output=outputfilename
```

The other parameters can be specified on the PyRAF command line or alternatively can be edited using the standard IRAF 'epar' mechanism before running the task. MultiDrizzle is designed to carry out the following steps, either in a single pass or alternatively by selecting various steps individually:

1. Staticmask - Identify negative bad pixels, based on examining all the images, and include them in the dq file
2. Skysub - Sky-subtract each frame
3. Driz_separate - Drizzle the input images onto separate, registered outputs (using shifts computed from the headers)
4. Median - Create a median image from the separate drizzled images
5. Blot - Blot the median image back to the original input frames
6. Driz_cr - Use each blotted image to create a derivative image, and compute CR masks
7. Driz_combine - Do the final drizzle combination

The parameters to each of these steps, as well as the specifics of their behavior, are described in more detail in the on-line documentation available within PyRAF for MultiDrizzle. There are also "startup" parameters which include whether or not to specify an output reference frame, or whether to use a shiftfile that may have been generated by Tweakshifts. Here we summarize the parameters for each of the seven steps mentioned above, pointing out some of the relevant issues that may need to be considered for various scientific applications.

3.1. Create the "Static" Mask of Negative Pixels

Parameters:

`static_sig`	=	4.0	Sigma value to use in flagging negative pixels

This step goes through each of the input images, calculates the r.m.s value for each chip, and identifies pixels that are below the median value by more than some number times the r.m.s. This is aimed at identifying pixels that may have high values in the dark frame that is subtracted during calibration, but may not necessarily have high values in the images, thus the subtraction gives them strongly negative values. Such pixels are not always flagged in the data quality array, hence this step allows them to be identified. Sometimes such pixels fall on bright objects so they would not be negative, but instead would be positive although lower than surrounding pixels. However, if the images are dithered then they should land on blank sky at least some of the time, in which case they will appear negative and will be flagged. The reason for identifying such pixels is to avoid problems later on when creating the median image.

3.2. Perform Sky Subtraction

Parameters:

`skywidth`	=	0.1	Bin width for sampling sky statistics (sigma)
`skystat`	=	'median|mode|mean'	Sky correction statistics parameter
`skylower`	=	INDEF	Lower limit of usable data for sky (in DN)
`skyupper`	=	INDEF	Upper limit of usable data for sky (in DN)
`skyclip`	=	5	Number of clipping iterations
`skylsigma`	=	4.0	Lower side clipping factor (in sigma)
`skyusigma`	=	4.0	Upper side clipping factor (in sigma)
`skyuser`	=	"	Header keyword containing sky value

This step calculates the sky value, using iterative sigma-clipping if specified. For instruments with multiple detectors, such as ACS/WFC and WFPC2, it calculates the sky separately for all the detectors, then chooses the lowest value and subtracts that from all the detectors. The reason for this is that the presence of bright sources on one detector can create a biased value of the sky measurement, thus the more accurate sky measurement is always taken to be the one that is lowest, since that is the least affected by any potential bright sources. If the 'workinplace' parameter is set to 'no' (the default), then MultiDrizzle will create copies of the input files and subtract the sky from those, leaving the original files unsubtracted. Setting this parameter to 'yes' will cause sky to be subtracted from the original input exposures.

3.3. Create Separate Drizzled Images

Parameters:

`driz_sep_outnx`	=		Output image x-size
`driz_sep_outny`	=		Output image y-size
`driz_sep_kernel`	=	'square|point|gaussian|turbo|...'	Drizzle kernel
`driz_sep_wt_scl`	=	'exptime|expsq'	Weighting factor
`driz_sep_scale`	=	'INDEF'	Output pixel size (arcsec)
`driz_sep_pixfrac`	=	1.0	Drop size, in input pixels
`driz_sep_rot`	=	INDEF	Output y-axis position angle
`driz_sep_fillval`	=	INDEF	Value for undefined pixels
`driz_sep_bits`	=	0.	Flag values considered good

This step drizzles the input images onto separate output images. By default it uses the *drizzle* 'turbo' kernel, and *drizzle* parameters of `pixfrac` = 1 and `scale` = 1. Specifying 'INDEF' for the scale means that it will set the output pixel size to the native pixel scale of

the camera, whatever that may be. These values can be changed depending on the scientific goals; for example, masks can be substantially improved by specifying a smaller value of `scale`, with the trade-off being larger images (their size increases as the inverse square of the value of `scale`), and increased computation time. The 'bits' parameter can be set to non-zero values if there are specific classes of pixel flag values that can be considered good. For example, the ACS data quality arrays contain large numbers of pixels with bit values of 32 and 64, which are often good in the images, thus setting 'driz_sep_bits' = 96 will allow both of these types of pixels to be treated as good, which means they will contribute to creating the median image in the next step. This may often be desirable if only a few exposures are being combined, where as many good pixels as possible are needed.

3.4. Create the Median Image

Parameters:

`median_newmasks`	=	yes	Create new masks?
`combine_type`	=	'minmed|average|median'	Type of combine operation
`combine_nsigma`	=	4 3	Significance for min. vs median
`combine_nlow`	=	0	Number of low pixels to reject
`combine_nhigh`	=	1	Number of high pixels to reject
`combine_lthresh`	=	INDEF	Low threshold for clipping
`combine_hthresh`	=	INDEF	High threshold for clipping
`combine_grow`	=	1.0	Radius for neighbor rejection

This creates a median image from the separate drizzled input images, allowing a variety of combination and rejection schemes. If `combine_type` is set to 'median' or 'average', then the routine behaves similarly to the IRAF task *imcombine*, using the values of `combine_nlow` and `combine_nhigh` (the number of low and high pixels to reject) and `combine_grow`, the amount by which flagged pixels can grow. If `median_newmasks` = 'yes', then pixels are flagged using the static bad pixel masks. If this parameter is 'no' then this step will simply use whatever masks are specified in the 'BPM' header keyword of each image (which could be created by the user). In general, however, it is recommended to use the static bad pixel masks that are generated by default.

If `combine_type` is set to 'minmed', then this step will use a slightly more sophisticated algorithm to create a cleaner combined image. The basic concept in this case is that each pixel in the output combined image will be either the median or the minimum of the input pixel values, depending on whether the median is above the minimum by more than a certain number of sigma. An estimate of the 'true' counts is obtained from the median image (after rejecting the highest-valued pixel), while the minimum is actually the minimum unmasked ('good') pixel. This algorithm is designed to perform optimally in the case of combining only a few images (3 or 4), where triple-incidence cosmic rays often pose a serious problem for more simplified median combination strategies. It performs the following steps:

1. Create median image, rejecting the highest pixel and applying masks

2. Use this median to estimate the true counts, and thus derive an r.m.s.

3. If the median is above the lowest pixel value by less than the first value mentioned in `combine_nsigma`, then use the median value, otherwise use the lowest value.

If `combine_grow` > 0, repeat the above 3 steps for all pixels around those that have already been chosen as the minimum, this time using a lower significance threshold specified as the second value in `combine_nsigma`. This is very successful at flagging the lower-S/N 'halos' around bright cosmic rays that were flagged in the first pass.

3.5. Blot Back the Median to the Frame of the Original Images

Parameters:

`blot_interp`	=	'poly5|poly3|nearest|linear|sinc'	Interpolant
`blot_sinscl`	=	1.0	Scale for sinc interpolation kernel

This takes the median image and uses *blot* to apply the geometric distortion and transform it back to the reference frame of each of the original individual input images, in preparation for the subsequent step of cosmic-ray rejection. Since *blot* uses interpolation to map the pixels back to the input detector frame, it is possible to choose the type of interpolant, as well as varying the scale if the sinc interpolant is used. Generally the default interpolant of 'poly5' is sufficient, and access to the other interpolants is provided purely on the basis of flexibility.

3.6. Create Cosmic Ray Masks

Parameters:

`driz_cr_corr`	=	no	Create CR-cleaned _cor file and _crmask file?
`driz_cr_snr`	=	'3.5 3.0'	driz_cr.SNR parameter
`driz_cr_grow`	=	1	Radius to grow around flagged pixels
`driz_cte_grow`	=	0	Length to grow along CTE direction
`driz_cr_scale`	=	'1.2 0.7'	driz_cr.scale parameter

This uses the original input images, the blotted images, and the derivative of the blotted images (created using the *deriv* task) to create cosmic ray masks (using the *driz_cr* task), stored as separate files, which can later be combined with other masks. This step can also create a '_cor' image, where bad pixels are replaced with pixels from the blotted median image. These relatively clean '_cor' images can also be used to determine shifts. The cosmic ray mask can be 'grown' by a certain pixel width, which can help eliminate faint halos around cosmic rays. In addition, it can be grown specifically along the CTE direction (typically $\sim 10 - 20$ pixels) which is useful in datasets where cosmic rays have CTE trails.

3.7. Perform Final Drizzle Combination

Parameters:

`final_wht_type`	=	'EXP|ERR|IVM'	Type of weight mask
`final_outnx`	=		Output image x-size
`final_outny`	=		Output image y-size
`final_kernel`	=	'square|point|gaussian|turbo|...'	Drizzle kernel
`final_wt_scale`	=	'exptime|expsq'	Weighting factor
`final_scale`	=	INDEF	Size of output pixels
`final_pixfrac`	=	1.0	Size of 'drop'
`final_rot`	=	0.0	Rotation (anticlockwise)
`final_fillval`	=	INDEF	Value for undefined pixels
`final_bits`	=	0.	Flag values considered good

This takes the original input images, together with the final cosmic ray masks, and drizzles them all onto a single output image. The standard *drizzle* parameters of `kernel`, `scale`, `pixfrac` and `rot` can be specified for this step. By default the pixel scale of the output image is 1, but feel free to experiment with other options (e.g. when combining at least 4 sub-pixel dithered images, `scale` = 0.5 and `pixfrac` = 0.7 can yield a sharper output PSF). By default, it creates an output weight image that is simply the effective exposure time ('EXP'); it is also possible to request it to create an inverse variance weight mask which takes all the information in the '[ERR]' extension in ACS FLT files. Users can also create their own inverse variance mask, in which case 'IVM' is selected and the mask filenames are provided with the input exposures in an ASCII file that is given to the 'input' parameter.

4. Tweakshifts - Real-Time Shift Refinement

Although the relative astrometric information in the image headers is generally reliable to better than 0.1 pixel for data obtained during the same visit and with the same guide stars, there are often cases where the shifts need to be improved. These include the following situations:

- Visits that had guide star problems, which can produce drifts from one exposure to the next depending on how the telescope was tracking. Sometimes these drifts are on the order of less than a few pixels between exposures, which generally means that it may still be useful to combine them. However, if the drift rate is much larger then this will also be noticeable on the exposures themselves and they may need to be excluded from the combination.

- Science programs that require much better relative positioning between exposures than the nominal $\sim 0.05 - 0.1$ pixel accuracy achievable from the header astrometry. This can be required for very precise stellar photometric or astrometric measurements, for example. Such data can be amenable to much more accurate relative registration if they contain more than a few thousand stars, since the aggregate uncertainty in the cumulative shift measured from all the stars can be reduced from the nominal ~ 0.1 pixel centroiding accuracy for a single star down to $\sim 0.001 - 0.003$ pixels for the net shift measurements if several thousand stars are used.

- Combination or comparison of data from different visits, where different guide stars may have been used, in which case the header astrometry between the visits may differ by as much as a few arcseconds. In these cases, if there are enough objects in common between the overlapping exposures then they can be used to directly align the images. On the other hand, if the exposures are part of a mosaic pattern with minimal overlap, then it may be better to register the images separately to a high-quality external catalog of the entire field, which will then provide sufficiently accurate relative registration between them. This is the strategy employed for several of the large-scale imaging programs on *HST*.

- Exposures of moving targets within our solar system, for which the alignment may need to be improved if MultiDrizzle does not recognize the ephemeris information, or if the ephemeris information is not completely accurate, in which case the target may be at an unexpected location on the detector.

The Tweakshifts script has therefore been developed in order to provide a means to solve for these shifts, in a relatively automatic way. It can be run in several different modes, depending on the scientific requirements:

- Refine the relative shifts of a set of exposures, either from the same visit or from different visits (assuming sufficient overlap).

- Register an image to a pre-existing output reference image.

- Register an image to a pre-existing catalog.

When using Tweakshifts to align images to one another, the observer can select either cataloging or cross-correlation as the technique to use in solving for shifts. Cataloging can be carried out using either SExtractor (Bertin & Arnouts 1996) or DAOPHOT (Stetson 1987), and is generally appropriate when the images contain a large number of discrete sources, such as star clusters or sparse extragalactic fields. When the images contain large, extended emission, such as galactic nebulae, large external galaxies or even solar system targets, then it is often more appropriate to use cross-correlation, since a lot of the signal

in the images may be contained in large-scale diffuse structure that may prevent effective cataloging but can provide a robust cross-correlation measurement.

Tweakshifts provides a significant amount of access to the parameters for the various techniques available to solve for shifts. All the SExtractor parameters are available, as are the relevant parameters for DAOFIND. When an input catalog is provided, Tweakshifts will use the R.A. and Dec. columns as input (as specified by parameters), and will use transformations to the distorted frame of an individual exposure in order to compute the required shifts.

The resulting shifts are written out to an ASCII text file for easy manipulation, and can also be stored in an association table if that was provided as input to Tweakshifts. Thus, when MultiDrizzle is run, it can be given the shiftfile as input (or the modified association table), and will correctly apply the shifts to each exposure in order to improve their registration.

The first public release of Tweakshifts in PyRAF STSDAS V3.4 (1 November 2005) to the community is considered a prototype, but has sufficient flexibility and robustness that it should be directly useable for a significant majority of scientific applications that require it. Future versions will include more functionality and other improvements in response to feedback from the community.

5. The Future - Virtual Observatory and Hubble Legacy Archive

Since the fundamental design of both MultiDrizzle and Tweakshifts is aimed at allowing them to be run autonomously on a large variety of datasets, this makes them amenable to incorporation into the *HST* Archive Pipelines that automatically process *HST* data. MultiDrizzle was incorporated into the *HST* on-the-fly-reprocessing (OTFR) pipeline for ACS in September 2004, and since then it has been delivering automatically cleaned, drizzled, combined images for all ACS associations that contain multiple exposures of a given target.

As a result, it is conceivable that the use of MultiDrizzle can be extended to provide products that are suitable for the Virtual Observatory and Hubble Legacy Archive (VO/HLA), which will be aimed at delivering clean, combined, geometrically rectified images for a variety of different purposes. For example, enabling scientific work to be carried out on multiple visits of the same target, either by combining them or by enabling time-variable phenomena to be searched for in different epochs, will require automatically cleaned images of each epoch, registered onto a common grid.

Since Tweakshifts is amenable to being run autonomously to determine the relative registration between datasets, this means that it could in principle be used in the VO/HLA context to refine the shifts between any set of images specified by the observer. Another use for Tweakshifts would be to register images to an existing, pre-defined catalog, for example the Guide Star Catalog II (GSC-II; McLean et al. 2004; McLean 2006), which will be used in *HST* operations from Cycle 15 onward and has much improved astrometry over the previous system. A demonstration of the potential of this technique has recently been carried out for a significant number of images from the first year of ACS operations (Koekemoer et al. 2005b, and these proceedings). Current work is aimed at investigating ways in which Tweakshifts can be made more robust, as well as incorporating new algorithms that may make it applicable to an even wider range of datasets.

Acknowledgments. We are pleased to acknowledge very valuable contributions from a large number of people who have contributed ideas or feedback, including Eddie Bergeron, Tom Brown, Ivo Busko, Stefano Casertano, Colin Cox, James Davies, Linda Dressel, Harry Ferguson, Shireen Gonzaga, Perry Greenfield, Dave Grumm, Robert Jedrzejewski, Diane Karakla, Ray Lucas, Jennifer Mack, Jesus Maiz-Appelaniz, Max Mutchler, Vera Platais-Kozhurina, Adam Riess, Marin Richardson, Ed Smith, Megan Sosey, as well as many others.

References

Bertin, E. & Arnouts, S. 1996, A&AS, 117, 393

Fruchter, A. S. & Hook, R. N. 2002, PASP, 114, 144

Gilliland, R. 2005, "Guiding Errors in 3-Gyro: Experience from WF/PC, WFPC2, STIS, NICMOS, and ACS", *Instrument Science Report* TEL 2005-02, (Baltimore: STScI), available through http://www.stsci.edu/hst/observatory/documents/isrs

Koekemoer, A. M. et al. 2002a, "The *HST* Dither Handbook V2.0" (Baltimore: STScI)

Koekemoer, A. M., Fruchter, A. S., Hook, R. N., & Hack, W. 2002b,2003, in *Proc. 2002 HST Calibration Workshop*, ed. S. Arribas, A. Koekemoer, & B. Whitmore (Baltimore: STScI), p. 337

Koekemoer, A. M., Kozhurina-Platais, V., Riess, A., Sirianni, M., Biretta, J., & Pavlovsky, C. 2005a, "Two-Gyro Pointing Stability of *HST* Measured with ACS", *Instrument Science Report* ACS 2005-07, (Baltimore: STScI), available through http://www.stsci.edu/hst/acs

Koekemoer, A. M., McLean, B., McMaster, M., & Jenkner, H. 2005b, "Demonstration of a Significant Improvement in the Astrometric Accuracy of *HST* Data", *Instrument Science Report* ACS 2005-06, (Baltimore: STScI)

McLean, B., 2006, The 2005 *HST Calibration Workshop.* Eds. A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel, this volume, 413

McLean, B., Lattanzi, M., Greene, G., Loomis, C. & Meakes, M. 2004, SpaceOps 2004 conference proceedings

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

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Band-limited Imaging with Undersampled Detectors

Andrew S. Fruchter

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. Over the past decade "Drizzle" has become a de facto standard for the combination of HST images. However, the drizzle algorithm was developed with small, faint, partially-resolved sources in mind, and is not the best possible algorithm for high signal-to-noise unresolved objects. Here, a new method for creating band-limited images from undersampled data is presented. The method uses a drizzled image as a first order approximation and then rapidly converges toward a band-limited image which fits the data given the statistical weighting provided by the drizzled image. The method, named iDrizzle, for iterative Drizzle, eliminates the small high-frequency artifacts that can be introduced by drizzling. The method works well in the presence of geometric distortion, and can easily handle cosmic rays, bad pixels, or other missing data. It can combine images taken with random dithers, though the number of dithers required to obtain a good final image depends in part on the quality of the dither placements.

1. Introduction

The competition between the desire to have a wide field-of-view and the limitations on the number of available pixels – caused by a desire to reduce cost and complexity or to lessen read-noise – means that astronomical detectors are often undersampled. In order to fully sample an image with an undersampled detector one must dither and combine multiple images. However, distortions in the field-of-view may make it impossible to perform shifts that equally well sample different parts of the detector. In practice then the combined pixels from dithering of astronomical detectors often produce irregular sampling of the image plane.

In order to combine the irregularly sampled data from the Hubble Deep Field *HDF* (Williams et al. 1996), a new image algorithm, Drizzle (Fruchter and Hook 2002) was developed. Drizzle combines dithered images in a statistically optimal fashion. However, as can be seen in Figure 1, Drizzle adds small high-frequency artifacts to the image. On scales larger than an original pixel, these rapidly average out. Thus for the prime purpose of the *HDF*, the study of small faint galaxies, Drizzle is an excellent algorithm. However, for the analysis of high signal-to-noise images of point sources, or other cases where preservation of the true point spread function (PSF) was essential, one might prefer an algorithm which more exactly reproduced the true band-limited image that is formed in the detector. While there is no known method to analytically reconstruct an irregularly-sampled band-limited image, here we present an iterative method which performs. In essence it is an iterative approximation to the band-limited image, using Drizzle to ensure that the data are combined using the full statistical power of the data. In the absence of noise this method converges directly to the true image. In the presence of noise a small increase in the statistical noise is produced, but the systematic high-frequency noise introduced by Drizzle is removed.

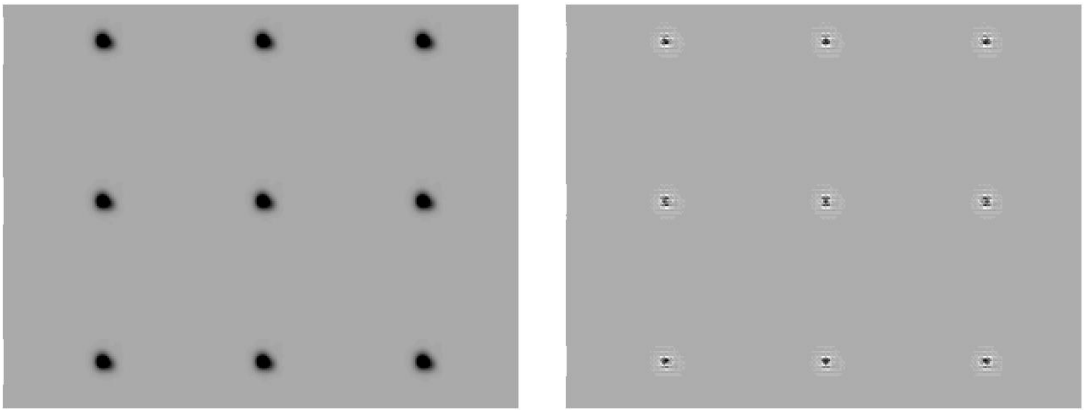

Figure 1: On the left a series of synthetic ACS PSFs. At the right, a drizzled approximation of the image subtracted from the original. The largest residuals are slightly greater than 10% of peak. Drizzle smoothes the PSF and adds high frequency noise. The smoothing is reproducible; thus if a PSF slightly larger than the original is acceptable, the smoothing is not a significant issue. Similarly, if one is measuring properties of the image on scales larger than a couple of original pixels, the high frequency noise largely averages out. Thus Drizzle is well suited for aperture photometry with apertures larger than 2 original pixels, or galaxy photometry. If one wished to reconstruct a true instrumental PSF, or if one wishes to do photometry through PSF fitting, another method might be preferred.

2. The Method

The spatial frequencies in an image are limited by the optics of the telescope to be $\leq D/\lambda_s$, where D is the (maximum) aperture of the telescope and λ_s is the shortest wavelength in the passband. Thus the spatial frequencies in an images are band-limited. One might imagine then that one could take the total set of irregularly sampled data which everywhere locally meets the Nyquist criterion and do a Fourier transform, remove any frequencies above the cutoff frequency, and Fourier transform back to get the true band-limited image. Unfortunately, a direct Fourier transform of an irregularly-sampled data set throws a great deal of power out of the original passband, and thus this method fails terribly. Surprisingly, it is more effective to first put the data onto a regular Nyquist grid by simply taking the value of the nearest neighbor before doing the Fourier transform. The inverted, band-limited function turns out to be a much truer approximation than the direct transform case. This approximation is know as the Voronoi approximation, and more information on it, and the ideas discussed in this paragraph can be found in the tutorial by Tobias Werther (2006).

Now the Voronoi approximation is a band-limited function and thus can be sinc interpolated to the irregular grid of the data. One can therefore subtract the Voronoi approximation from the original function at all of the data points. Furthermore, this smaller difference function is itself a band-limited function, so one can repeat the process and get a further refined approximation to the underlying band-limited function. This procedure is known to converge geometrically (Werther 2006).

The nearest neighbor approximation is, however, far from ideal for astronomical imaging. Only one data sample is used at any point on the regular grid, even if several nearby data samples could provide information, and there is no means to weight the data according to its statistical significance. Therefore in the proposed method the nearest neighbor approximation is replaced by Drizzle in each iteration. While Drizzle introduces small artifacts (as does the nearest neighbor) the iterative comparison with the original data serves to remove these, as can be seen in Figure 2.

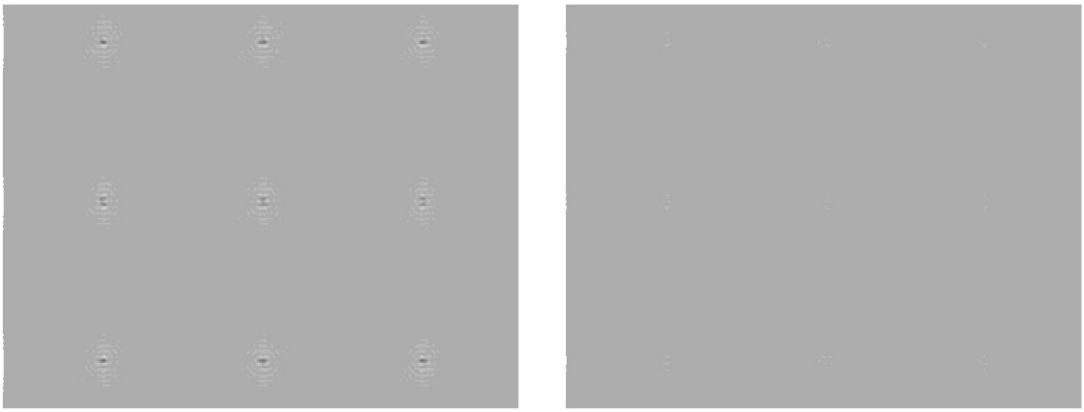

Figure 2: On the left, the residual difference between the true ACS PSF (see Figure 1) and the iterative approximation after the first iteration. On the right the same difference after several more iterations. The rapid reduction of the small artifacts of the Drizzle algorithm is evident. Twelve simulated noiseless ACS images with random pointings were combined.

It is noise which requires the use of Drizzle, rather than the nearest neighbor approximation, but in Figure 2, there is no noise. On the left-hand-side of Figure 3 we therefore show the same subtraction performed between noiseless ACS PSFs and a combination of twelve simulated ACS images of the stellar field, with each star near saturation in a 1200s exposure, with appropriate Poisson and read noise added. This tests the method in a situation of extremely high signal-to-noise. The residuals are close to that expected from noise statistics, and the peak errors are reduced from the Drizzled subtraction by a factor of ~ 20. However, the introduction of noise greatly slows convergence – twenty-four iterations were used to produced the output shown here.

On the right-hand-side of Figure 3 is a central region of the Hubble Ultra-Deep Field (Beckwith et al. 2006). The bright star in this image is near saturation in each of the twelve 1200s individual exposures combined with the new method to form the final image. When this image is mapped back onto the individual input exposures and subtracted, residuals of approximately 2% peak are found under the stellar image. These small but measurable errors are most likely caused by temporal variations in the PSF caused by variation in the insolation of the telescope as it orbits the earth.

3. Noise

3.1. Noise Amplification

Drizzle places the output image exactly where it was observed. But the average weight of an output pixel will not necessarily fall at the center of that pixel, and thus there is a jitter between the represented and effective position of a pixel. Furthermore the peak of a drizzled PSF will never be greater than the greatest value in the appropriate region of the input images. By contrast iDrizzle attempts to predict the true value of the image at the center of the output pixel, and thus the peak of a PSF will often be brighter than any value of the input images in the appropriate region. iDrizzle essentially uses estimates of the derivatives of the data to extrapolate to a position (the center of the output pixel) which is not necessarily exactly sampled in any of the input images. This will produce some noise amplification, which will vary with the quality of the dithering. In typical tests performed on this method the noise amplification has been in the region of 10%.

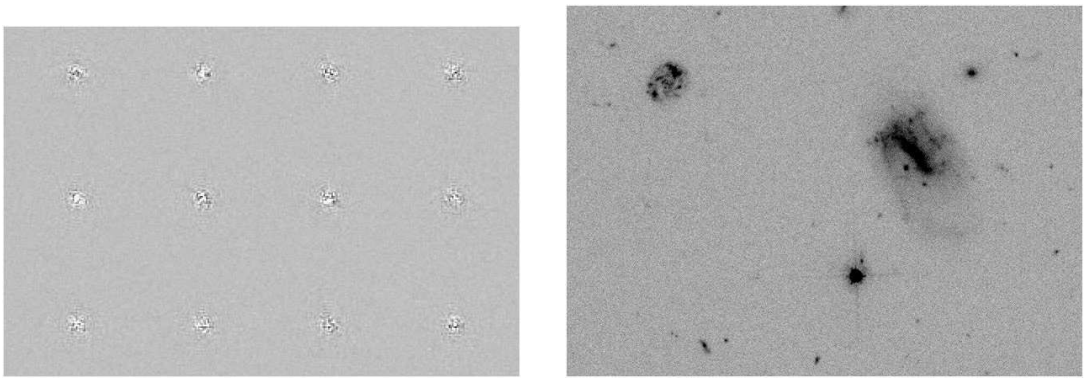

Figure 3: On the left, the same subtraction of the combination of twelve simulated ACS images from the true PSF as seen in Figure 2 but with Poisson and read noise incorporated. On the right, the combination of twelve images from the Hubble Ultra-Deep Field using the new method.

3.2. Correlated Noise

Images produced by Drizzle show correlated noise. Part of this correlation is caused by the drizzling process itself – a non-zero value of `pixfrac` causes Drizzle to place a given input pixel value down on a region of size $p \times s$, where p is the value of `pixfrac` and s is the size of an input pixel. As the iteration proceeds, iDrizzle effectively forces p to zero. But this does not entirely remove correlated noise.

To see why correlated noise remains consider the following situation: the center of an input pixel falls directly on the boundary between two output pixels. Because the image is band-limited (and presumably the output image uses a sampling at least as fine as Nyquist) this input value must affect the predicted value of the two pixels on either side of it. Thus any noise in this input pixel will affect the noise in these two pixels, and thus the noise in these pixels will be correlated. Thus this method reduces but does not eliminate correlated noise.

4. Final Comments

In this short paper, a new method for the combination of dithered astronomical images has been introduced. The method has the ability to handle shifts, distortions and missing data, and converges rapidly to an accurate representation of the underlying image. It requires, however, that the combined images Nyquist sample the image. In the case of random placement of images this may take a minimum of order eight and twelve images for the ACS and WFPC2, respectively, on *HST*.

This article is intended as an outline for iDrizzle. Details – such as apodization of the image and the cut in the Fourier plane – are defered to later and longer presentations. Here, the essence of the method, and its benefits and limitations are presented.

References

Beckwith, S. V. W. et al. 2006, AJ, submitted

Fruchter, A. S. & Hook, R. N. 2002, PASP, 114, 144

Werther, T. 2006, http://www.math.ucdavis.edu/štrohmer/research/sampling/irsampl.html

Williams, R. E. et al. 1996, AJ, 112, 1335

The 2005 HST Calibration Workshop
Space Telescope Science Institute, 2005
A. M. Koekemoer, P. Goudfrooij, and L. L. Dressel, eds.

Where Will PyRAF Lead Us? The Future of Data Analysis Software at STScI

P. Greenfield

R. L. White

Space Telescope Science Institute, Baltimore, MD 21218

Abstract. PyRAF, which allow the running and scripting of IRAF tasks from Python, was just the first step in STScI's migration to a more flexible and productive data analysis environment. The ultimate goal is to combine the best of the IRAF and IDL approaches to data analysis software into one environment, one where astronomers can find it easy to write software for, and software developers have sufficient tools and capabilities to handle writing general applications and pipelines. We are in the process of building the infrastructure to support these goals, and are well along in doing so. This paper will discuss the tools already developed, including PyRAF, numarray, PyFITS, numdisplay, and matplotlib, as well as the tools that are in development or are planned. Finally, it will review the applications that have been or are being developed and the implications for the user community of our move in this direction.

1. Introduction

The PyRAF project was started to achieve both short-term and long-term goals. The short term goal was to provide a more robust and flexible scripting environment for IRAF tasks than the IRAF CL provided. At the time we were trying to maintain drizzle-related scripts that had gotten quite lengthy and were proving difficult to maintain. The absence of any error handling mechanism and easy identification of where errors occurred slowed progress considerably. An alternate CL looked like a promising way to overcome some of the existing limitations.

There were longer-range goals as well. These were not part of the initial development yet provided additional motivation for starting work on PyRAF. One was that it would enable easier integration of software not written under the IRAF environment with IRAF tasks. IRAF presents a comparatively closed software environment in that it is generally difficult to link other software tasks and libraries with IRAF tasks. While it is possible to run other software as "foreign" tasks from the IRAF CL, it still means invoking many separate processes and is a comparatively awkward way of integrating software. With a new CL, we could allow much easier integration of other libraries and be free of the restrictions that IRAF places on such software.

Finally, A new CL would enable the possibility of providing IDL-like data manipulation facilities, something that had long been desired by many users.

2. Why Python?

We gave no thought to developing a new, custom scripting language. Developing and sustaining a good scripting language is a major effort. Past experience indicates that as-

tronomy is not a large enough field to warrant such an effort. Many scripting languages already existed and had sizable development communities.

In effect, one such 'scripting' language was already in wide use in astronomy, namely IDL. Granted, it isn't generally regarded as such, and although it isn't as general purpose as many of these, it shares at least a few characteristics with the better known ones to be viewed as one. Given its wide prevalence in astronomy, it was an obvious choice for trying as an alternate CL for IRAF. It's use would automatically satisfy the second of our long-term goals. Unfortunately, it lacked the necessary features to run IRAF tasks at the level of control we desired. This was the primary reason for rejecting it, though there were other important, but more secondary reasons. These included its cost, poor support for easily distributed C or Fortran extensions, and that it was not as general-purpose as other scripting languages, both in its language and library capabilities.

We deemed it important that any selected scripting language be popular enough, and have a bright enough future that it would be expected to be around many more years. At the time, there were three popular open source scripting languages: Perl, Tcl, and Python. The first, while very powerful and the most popular, was (and still is) considered difficult to learn, quirky, and led to code that was difficult to read, understand, and maintain. It generally is difficult to write large, maintainable programs in Perl. Tcl, on the other hand, has poor support for numerical types that rendered it unsuitable for astronomical contexts.

That left Python. It appeared to have all the necessary qualities including:

- Free, Open Source
- Very strong (deep and broad) user and development community
  - Still growing
  - Long lifetime expected
- General purpose
- Very portable
- Very extensible (in Python and with compiled code)
- Comparatively easy to learn
  - Yet very powerful
- Scales well

It was the clear choice at the time, and continues to be (http://www.python.org).

Initially, our goals were modest. We were not seeking to support IRAF graphics nor IRAF CL syntax. As the project progressed, we found that both these capabilities were within reach without enormous effort.

3. PyRAF Capabilities

Work on PyRAF began in 1998, and by mid-1999 it was being used internally at STScI. The first public beta release occurred in spring 2000. PyRAF provided a number of capabilities that were not present in IRAF at the time (a form of error handling and error traceback, as well as more conventional command line recall, have recently been added to IRAF, and are available in a beta release). The most notable aspect of PyRAF is that it doesn't require any changes to IRAF itself. It coexists with the existing IRAF CL; one can run one or the other once PyRAF is installed. PyRAF provides enhancements for interactive users and script writers.

For interactive users, the enhancements include:

- IRAF CL syntax accepted (as well as Python syntax)
- Ability to run most CL scripts
- Simple, up-arrow command line recall
- Tab completion of file names and commands
- Resizable GUI Graphics windows with plot recall
- GUI epar editor

For scripting, the following is possible:

- Full Python exception handling[1]
- Traceback information on error location
- Capture of task standard output to Python list structure

Not all aspects of the IRAF CL are currently supported. None of the current omissions are intrinsically not possible. Many are planned to be added in future enhancements (all those marked with *). Some arise from the fact that the IRAF CL behavior is not well documented, and arise from user reports on discrepancies in behavior. These are noted and fixed as needed (more minor discrepancies remain and will be fixed in future releases). Some discrepancies will never be corrected as they highlight IRAF features that are best not emulated since they allow poor programming practices. Examples of these sorts of features include the ability to minimum match declared procedure variables (one may use x in place of xcoord if nothing else starts with x) or change variables outside of the scope of the procedure. In those cases we strongly recommend that IRAF CL scripts that use them be modified not to depend on such features (this is always possible and generally not difficult to do). The following list gives the major differences that currently remain.

- GOTOs not supported (next release will support forward GOTOs; already implemented).
- Background tasks*
- Package unloading*
- Text-based epar (for remote terminal use)*
- Graphics redirection*
- Emulation of new IRAF CL error handling*

In addition to the currently marked differences between the IRAF CL and PyRAF, we plan many enhancements for PyRAF in the next year. Although we are no longer writing much software that requires PyRAF functionally (i.e., our new software no longer is written within the IRAF environment), it is still very important for our user community since it allows our new software to be run in an environment that is familiar to IRAF users and provides a nearly identical user interface. It also will play a key role in the Gemini calibration pipelines.

Further information may be found at:
http://www.stsci.edu/resources/software_hardware/pyraf

[1]Note that the new error handling capability provided in IRAF ECL beta release does not provide exception handling in the usual sense. It executes the entire block of code and then jumps to the error handler at the end regardless of where the error occurs. Most exception handlers jump immediately on the error.

4. Python's Larger Role

It is becoming well recognized by many software developers in astronomy (if not yet astronomers) that Python is a very good choice as a scripting language. Indeed, many of the existing or new astronomical software now sport, or will soon, some sort of Python interface. For example:

- Midas: PyMidas
- AIPS: Parseltongue
- ALMA
- CIAO/s-lang: PySL

The availability of Python interfaces to these systems will make it far easier to write scripts that combine capabilities from different systems.

Nevertheless, it is still a common viewpoint that Python is primarily only a scripting language and that it doesn't have a larger role to play in writing astronomical data analysis and reduction software. We have a more ambitious view. Our success in developing PyRAF led us to believe that it was a powerful language with the potential for being a primary, rather than secondary, language to develop applications in, provided the necessary tools were developed.

This was inspired, in part, by our previous experience with IDL that convinced us that it was possible to use an interpreted language to develop many useful applications. The presence of array manipulation facilities allows many computationally-intensive algorithms to be efficiently implemented in the interpreted language. While it is admittedly true that there are many scientific computation problems that are not well suited for array manipulations (e.g., fluid dynamics computations), our experience (and evidently the experience of many astronomers as well, given IDL's popularity in astronomy) was that such an approach was well suited to many data analysis and reduction problems.

No single language is suitable for all problems. Python is no exception. Yet, it is more sensible in our view to develop as much as possible in Python because of its highly productive nature than to take the more conventional approach of doing most development in a compiled language such as Fortran (any variant), C, C++, or Java, and then use a scripting language. The differences in productivity between the two approaches are startling, both in lines of code required and development time. Typically one sees increases in productivity ranging from factors 5 to 10 when switching from one of the traditional compiled languages to Python (much the same improvements have been seen by IDL users as well). Since Python provides very good interfaces to C (and thus to Fortran and C++), any part of the code that is too slow can instead call compiled code. Even with heavily computational code, generally speaking, most of the time spent in a program is usually confined to 10-20% of the code. Even if that code cannot be efficiently written in Python, it can be rewritten in C or Fortran. In other words: minimize the use of the less productive language.

An important side benefit of this approach is that it makes the code developed by professional software developers much more accessible to astronomers, and makes it easier for astronomers to contribute algorithms and utilities to a common software base. Currently, the two communities have little overlap in the software they develop; ending this divide will have many benefits.

5. Enabling Python as a Data Reduction Language

Were the necessary facilities already available in Python? In part. There was an array package, Numeric (aka numpy) already available. However, it had a number of features

and characteristics that made it less desirable for dealing with the large data sets we see in astronomy. It did not have the facilities for making efficient use of memory, nor could it support an efficient use of table data structures. We also desired to use memory mapping to minimize memory usage. For these and other reasons we decided on developing a new array package, numarray (http://www.stsci.edu/resources/software_hardware/numarray).

Another essential library needed to support the use of Python for applications was a FITS module. This was initially developed by Paul Barrett and then adapted to use numarray and provide table access. Despite being written in pure Python, it is significantly faster for reading large files than the standard CFITSIO library, particularly for very large files. (http://www.stsci.edu/resources/software_hardware/pyfits)

Before Python could be considered a serious interactive environment for data analysis, it had to have a capable way of visualizing data. We developed numdisplay to display images to DS9 and ximtool. But for a long time, one of the biggest holes was the absence of a good plotting package that met all our needs (which included being free, open source, portable to all popular platforms, GUI-toolkit independent, image capable, and having high quality hardcopy). This was finally filled with the adoption of the plotting package matplotlib developed by John Hunter at the University of Chicago (http://matplotlib.sourceforge.net/).

Besides extending the capabilities of these existing tools, efforts are beginning on addressing more astronomically-specific libraries that contain utilities similar to those found in the IDL astron library to handle coordinate transformations, WCS coordinate systems, and photometry. Most of this work is taking place on a public repository to encourage contributions and enhancements from the general community.
(http://www.scipy.org/wikis/topical_software/AstroLib)

Finally, there is currently an effort underway to unify the best aspects of the Numeric and numarray array packages into one called scipy_core (which will be as easy to install as as Numeric or numarray). We have done initial testing of scipy_core and it appears to satisfy all of the functional requirements we have. The effort to port our libraries and applications has begun. While scipy_core does not have an identical interface to numarray, it is much closer to numarray than Numeric in that regard. We will endeavor to make our software coexist with both during some reasonable transition period.

6. STScI Python Applications and Documentation

Most of STScI data analysis and reduction applications development is currently being done in Python. The following list summarizes those that have already been distributed or are completed and soon will be distributed. Most of these do not require IRAF at all (but many can be run as IRAF tasks in PyRAF using the standard IRAF CL syntax and epar interface). Most can be run on Microsoft Windows (though PyRAF cannot be because of its current dependence on IRAF).

- PyDrizzle/Multidrizzle: Drizzles a set of HST images onto a common image. Multidrizzle can remove cosmic rays even without any common pointings.
  (http://stsdas.stsci.edu/multidrizzle/)
- CALCOS: The calibration pipeline for the Cosmic Origins Spectrograph instrument for HST.
- FITSDIFF: A utility for comparing FITS files for differences in the data or headers. (Distributed with PyFITS)
- Various STIS tools (Distributed with stsci_python/STSDAS):
  - stisnoise: removes fixed pattern noise

  - sshift: aligns dithered spectral images

- PySPV: Python interface to Specview (Java GUI application)
- NICMOS saaclean tool: removes effects of SAA persistence (distrbuted with stsci_python/STSDAS)
- TFIT (from GOODS project): Does photometry on low resolution images using high resolution images for prior information.
- WCS matching tool: A GUI application for deriving WCS information for press release images from science images.

We have also developed a tutorial for astronomers showing how to use Python for interactive data analysis: http://www.scipy.org/wikis/topical_software/Tutorial

7. The Future of Python at STScI

The more we use Python the more convinced we are that we made the right decision in 1998 to base our work on it. With the maturation of the tools for implementing data analysis algorithms directly in Python we anticipate a substantial increase in the number of astronomer-written scripts that are shared in the community. The STScI data analysis software group is committed to continuing support and development of our Python tools and applications, and we plan to help coordinate the distribution of user scripts and libraries as well.

Acknowledgments. The various projects described here have involved many contributions by several people:

- PyRAF: Richard L. White, Perry Greenfield, Michele De La Pena, Warren Hack, Philip Hodge
- numarray: Jay Todd Miller, Perry Greenfield, Jin Chung Hsu
- PyFITS: Paul Barrett, Jin Chung Hsu
- numdisplay: Warren Hack
- matplotlib: created and maintained by John Hunter at the University of Chicago with STScI contributions from Nadia Dencheva, Paul Barrett, Jay Todd Miller, Perry Greenfield
- Regression Test Framework: Howard Bushouse, Bernie Simon, Hemant Shukla
- PyDrizzle/Multidrizzle: Warren Hack, Christopher Hanley, David Grumm, Robert Jedrzejewski
- CALCOS: Philip Hodge
- FITSDIFF: Jin Chung Hsu
- STIS tools: Paul Barrett
- NICMOS SAAClean: Victoria Laidler
- PySPV: Ivo Busko
- TFIT: Victoria Laidler
- Public Release WCS tool: Philip Hodge

Author Index

Subject Index