

Cassini Imaging Science Subsystem (ISS) Data User's Guide

Ben Knowles

Cassini Imaging Central Laboratory for Operations (CICLOPS)

Space Science Institute
4750 Walnut St., Suite 205
Boulder, CO 80301

9/29/2016

This document has been reviewed for ITAR compliance and has been determined not to contain export controlled technical data.

Table of Contents

1. Introduction.....	5
2. ISS Instrument Description.....	9
2.1 Mission Overview and Imaging Science Objectives	9
2.2 Camera Pointing Design.....	12
2.3 Camera Parameterization Part 1: Camera Settings and Operation	16
Filters.....	17
Shutter.....	20
Detector	21
Compression.....	23
Image Event	25
2.4 Camera parameterization Part 2: The ISS Pre-commanding Tool (ISSPT)	27
3. The ISS Data Set	30
3.1 ISS Science Data: Summary and Search Methods	30
3.2 Introduction to the PDS ISS Data Archive	41
3.3 ISS VICAR Image Format	44
3.4 Reading ISS Images	48
3.5 PDS Label.....	48
3.6 ISS Data Quality and Completeness	49
Truncated Images	51
Lossy Compression Camera Bug Anomaly.....	52
Double Bit Error Anomaly	53
NAC Haze Anomaly of 2001	54
3.7 Archive Volume Organization and Format.....	54
Archive Volume DVD Format.....	54
Calibration Volume Organization.....	56
Document File Formats	57
3.8 The Index.tab File	59
3.9 Cartographic Maps	60
3.10 Accessing ISS Data Online	61
4. ISS Image Calibration	66
4.1 Introduction to ISS Calibration.....	66
4.2 Theoretical Basis of Radiometric Calibration.....	67
4.3 ISS Calibration Data	72

Optics.....	73
Filters.....	76
Shutter.....	79
Detector.....	81
Point-Spread Function.....	89
Compression.....	91
Absolute Calibration.....	92
Polarimetric Calibration.....	96
4.4 CISSCAL: The Cassini ISS Calibration Software Pipeline.....	101
Introduction.....	101
Setting up the Environment.....	102
Basic Layout.....	103
Default Options File.....	104
Pull-Down Menus.....	105
Calibration Options.....	112
LUT conversion.....	113
Bit-weight correction.....	113
Subtract bias/2-Hz noise.....	114
Subtract dark.....	116
A-B pixel pairs.....	118
Linearize.....	119
Flatfield.....	119
Convert DN to flux.....	120
Correction factors.....	121
Geometric correction.....	123
Running CISSCAL from the IDL Command Line.....	123
Make Polar Image.....	125
5. End-to-End Example: Absolute Flux Calibration (West et. al., 2010).....	127
5.1 Introduction and Setup.....	127
5.2 Assembling the Data Set.....	131
Comparison Spectra.....	131
Summary of Data.....	132
Downloading Images from the Planetary Image Atlas.....	135
5.3 Image Calibration and Data Analysis.....	137
NAC Calibration and Photometry Procedure.....	137

Combining NAC Results.....	141
WAC Calibration and Photometry Procedure	143
BOTSIM Analysis	143
Combining WAC Results.....	145
Residual QE Correction.....	146
5.4 Final Results.....	147
Appendix A: ISS Instrument Data.....	148
Appendix B: ISS Observation Descriptions	160
Saturn Atmosphere	160
Saturn Rings	161
Icy Satellites.....	165
Small Satellites (aka “Rocks”)	167
Enceladus Plume	167
Calibration.....	167
Appendix C: Rev/Orbit and Sequence Boundaries.....	168
Appendix D: Bibliography	177

1. Introduction

This document intends to provide a comprehensive introduction to the data produced by the Cassini Imaging Science Subsystem (ISS), with emphasis on methods for obtaining, searching and radiometrically calibrating ISS images for purposes of scientific research. ISS data is archived with the Planetary Data System (PDS) in collaboration with the Cassini Imaging Central Laboratory for Operations (CICLOPS). At the time of this revision the ISS archive contains hundreds of thousands of images split between 85 volumes, along with an additional 11 volumes worth of calibration files, software, and ground calibration images. New data volumes will continue to be released, one every three months, until the end of the Cassini mission in 2017.

Users of this guide are assumed to have prior knowledge of astronomical CCD image processing and calibration. Those who wish to calibrate the images using the standard Cassini ISS Calibration software suite CISSCAL included on the PDS ISS Calibration volume will require access to a computer with the Interactive Data Language (IDL) installed. Any coding examples within this text will likewise be provided in IDL. Usage and familiarity with basic Linux/Unix routines and command line navigation is also assumed.

This guide is organized by subject such that each successive section builds on context presented in the previous sections. Section 2 contains a thorough description of the ISS instrument and its parameterization, including useful background on the pointing design and camera commanding processes. Section 3 introduces the PDS ISS data archive, including image and image label formats and information about accessing and making use of the PDS archive volumes and metadata. Particular attention is paid to likely search strategies and methods. Section 4 describes the radiometric calibration process, from its theoretical basis to its algorithmic implementation within CISSCAL, as well as additional image processing techniques not contained in CISSCAL such as polarimetric calibration.

Finally, section 5 will provide end-to-end examples that will follow the entire data pipeline from initial database search through image calibration to published result.

As will be discussed in detail in Section 3, ISS metadata largely takes the form of “KEYWORD=VALUE” pairs which are contained in both the image headers and detached PDS label files. To ensure clarity and facilitate familiarity with ISS keywords for searching, this document will always refer to these keywords using their official full names, which are given in all caps, with underscores in place of spaces (e.g. TARGET_DISTANCE). As on the archive volumes, image and label filenames will always be given in all caps, while all other supporting files and directory paths will always be in lowercase (although note that on the DVD volumes, depending on how the archive volume filesystem is interpreted by the host computer, this naming convention may not be retained). Directory names will always terminate with a “/” character.

In addition to the keyword metadata contained in the image labels, most of which is created either at the time of exposure or during initial post-downlink processing, CICLOPS also provides a comprehensive set of target geometry keywords whose values are calculated during the image auto-navigation (“Autonav”) process. These keywords are accessible via the index.tab file in the index/ subdirectory of the ISS archive volumes.

Much of the information covered in this document has been covered previously in various ISS, PDS and Cassini project documentation, and users are strongly encouraged to consult these sources in addition to this Data User's Guide in order to obtain the most thorough understanding of the ISS instrument and its data. Table 1 provides a summary of the most useful sources and references, and where they can be found. As for the ISS data itself, images may be downloaded directly from PDS in the original archive volume format here:

<http://pds-imaging.jpl.nasa.gov/volumes/iss.html>

or else via the Planetary Image Atlas search tool found here:

<http://pds-imaging.jpl.nasa.gov/search/search.html>

Note that as of volume COISS_2071 (covering images taken after October 1, 2011) individual archive volumes are no longer necessarily sized to fit on a single DVD. This is primarily due to the inclusion of the cumulative index file, cumindex.tab, which contains metadata for all images up through the current volume, and can be quite large in size.

Topic	Document(s)	Location
Cassini Spacecraft, Mission Background	insthost.cat, mission.cat	catalog/
ISS Instrument	issna_inst.cat, isswa_inst.cat	catalog/
ISS Instrument (detail)	Porco et. al. (2004)	Space Sci. Rev. 115, 363–497
ISS Data (low-level)	tlmtab.fmt, vicar2.fmt, prefix*	label/
ISS Data (anomalous data)	errata.txt, dataset.cat	./, catalog/
ISS Data (image label format)	edrsis.pdf	document/
ISS Data archive (including file formats, volume organization and full PDS keyword list)	aareadme.txt, archsis.pdf	./, document/
Calibration (theoretical basis)	theoretical_basis.pdf	document/ (Calibration volume)
Calibration (ground-based analysis and data tables)	Ground Calibration Report	document/report/ (Calibration volume)
Calibration (in-flight strategy)	in_flight.pdf	document/ (Calibration volume)
Calibration (CISSCAL software)	cisscal_manual.pdf	document/ (Calibration volume)
Calibration (CISSCAL data files)	Contents of CALIB/ directory and associated info and label files	calib/ (Calibration volume)
Calibration (in-flight status, latest results)	West et. al. (2010)	Planet. and Space Sci. 58, 1475-1488
Scientific Publication Reference List	projref.cat, issref.cat	catalog/
ISS Archive contact individuals	person.cat	catalog/

Table 1: List of useful reference documents for ISS.

ISS image calibration does not require using the detached PDS labels, or even necessarily any

of the support documentation found on the data archive volumes. It does, however, require the calibration software and support files located on the dynamic volume of the ISS Calibration Archive, the latest version of which is available here:

http://pds-imaging.jpl.nasa.gov/data/cassini/cassini_orbiter/coiss_0011_v2.tar.gz

Acquiring this volume will be the logical first step for most ISS data users.

This ISS Data User's Guide was written by Ben Knowles, calibration and archiving engineer for CICLOPS, under the direction of Imaging Team leader Carolyn Porco and Deputy Imaging Team leader Bob West. Portions of Section 3 are paraphrased from existing ISS archive volume documentation written by members of CICLOPS, including the aareadme.txt and dataset.cat. In addition, portions of Section 4 are taken from the CISSCAL user manual, which this document replaces and supersedes.

2. ISS Instrument Description

2.1 Mission Overview and Imaging Science Objectives

The Cassini spacecraft was launched on October 15, 1997, armed with a suite of twelve scientific instruments designed for the extended exploration of the Saturnian planetary system. In addition to being the most sophisticated mission ever sent to the outer solar system in terms of its scientific capabilities, upon reaching its destination on June 30, 2004 it became the most distant orbital outpost humankind has yet established. Since that time, Cassini and its companion mission to Titan, Huygens, have successfully acquired and transmitted back to Earth a wealth of information, allowing for an unprecedented look at some of the solar system's most varied, dynamic, and once-mysterious landscapes.

The Imaging Science Subsystem (ISS) is Cassini's highest-resolution 2-D imaging instrument, and as the primary optical navigation tool on board, can reasonably be considered the spacecraft's eyes. The ISS consists of two separate cameras, the Narrow-Angle Camera (NAC) and Wide-Angle Camera (WAC), which are boresight-aligned and have fields-of-view of 0.35 degrees and 3.5 degrees, respectively. The cameras have been designed for maximum flexibility: their photometric and spectral sensitivity, linearity and dynamic range, optical resolving power, and wide variety of compression and other data collection modes can accommodate a vast array of targets and imaging situations. In terms of real-world resolving capabilities, the NAC can achieve pixel scales as small as a few tens of meters on targeted satellites, and as small as 1 km on Saturn's rings and atmosphere.

Between its launch and when it finally arrived at Saturn (via a dramatic and complicated maneuver known as Saturn Orbital Insertion, or SOI), the Cassini spacecraft made its way to the outer solar system with the help of several gravity assists, including relatively close encounters with Venus,

Earth, Jupiter, and a serendipitous encounter with the asteroid Masursky. With the exception of periodic star calibration observations, these are the only events during which the ISS cameras were used prior to Saturn approach. The Jupiter flyby, in particular, served as a trial run for many of the basic pointing modes and uplink processes that would eventually be used at Saturn.

The science goals for the nominal Cassini mission are many and varied, reflecting the huge diversity of target objects in the Saturnian system, and the unprecedented view afforded by the spacecraft's suite of instruments. Table 2 contains a brief summary of imaging science objectives, organized by target; a much more thorough account of ISS science objectives can be found in Porco et. al. (2004).

Saturn Rings	<ul style="list-style-type: none"> • ring structure (including ring edges, waves, perturbations, satellite orbital interactions, gaps, embedded moons, and ringlets) • particle properties and disk structure • spokes • diffuse rings (E, F and G rings)
Saturn Atmosphere	<ul style="list-style-type: none"> • spectrophotometric/polarimetric studies • atmospheric dynamics • auroras
Satellites (non-Titan)	<ul style="list-style-type: none"> • surface morphology, cartography and topography • tectonism, crater studies, geomorphology • albedo and color measurements • satellite orbital parameters, mass
Titan	<ul style="list-style-type: none"> • meteorology • spectrophotometric/polarimetric studies • surface features

Table 2: Summary of imaging science objectives for Cassini as described in Porco et. al. (2004).

The nominal mission for Cassini was to last for 75 Saturn orbits through 2008, and has since been extended twice. The first extension, which ran for two years and included Saturn equinox in the August of 2009, was dubbed the “Cassini Equinox Mission” (CEM), or more colloquially, XM (“extended mission”). The second extension is set to last through 2017, and is called the “Cassini

Solstice Mission” (CSM, or XXM/”extended extended mission”) for the solstice that will occur in May of that year.

Particular highlights from the nominal mission include: approach “movies” of Saturn, it's rings, and Titan; a close flyby of Phoebe on approach (the only close encounter with that satellite during the mission); high-resolution imaging of the rings at SOI; 45 targeted flybys of Titan, including 24 within 1000 km of the surface; targeted or other “close” flybys of Iapetus, Hyperion, Rhea, Dione, Tethys, Mimas and Enceladus; three extended periods of inclined and low-periapse orbits for high-resolution, high-latitude Saturn and ring imaging; and a dozen orbits with apoapses distant enough to allow for global coverage of Saturn's atmosphere.

The Cassini Equinox Mission featured continued close surveillance of Titan, Rhea and Dione, as well as multiple targeted flybys of Enceladus, with particular attention given to the south polar region to investigate the active “tiger stripes” and their associated geyser-like plume extending above the surface. But the primary focus of this portion of the mission was the rings, and the unique viewing geometry afforded during the months leading up to and following Saturn equinox, revealing topological features that cannot be seen at any other time. This portion of the mission is known as the Equinox Rings campaign.

The Solstice Mission gives Cassini the opportunity to monitor the Saturnian system over a complete seasonal period. Targeted flybys of Titan, Rhea, Dione, and Enceladus will continue, along with a close non-targeted flyby of Helene, and a relatively distant first flyby of Pallene, a satellite first discovered by Cassini in 2004. Towards the end of the mission, the science emphasis will shift back to Saturn and its rings, culminating in two weeks of very close orbits before the spacecraft's dramatic final plunge into Saturn's atmosphere.

2.2 Camera Pointing Design

Pointing of the ISS cameras and the Cassini spacecraft in general is accomplished through a complex iterative process involving input from all the instrument science teams. This “uplink” process begins with science planning and negotiations between instruments and individual Target Working Teams (TWTs) and Orbiter Science Teams (OSTs) and ends with the actual execution of camera commands by the spacecraft.

Cassini's movements are constrained by over 300 flight rules, some of which are obvious – no pointing of sensitive instruments towards the sun, for example – and some of which are less so. To facilitate compliance with these constraints and reduce operations costs associated with daily spacecraft communication, the Cassini mission is split into extended periods of time – typically 5-6 weeks for the nominal mission but longer in the CSM - called “sequences.” Each sequence requires the generation of a file containing the full set of commands to be executed by the spacecraft during that period, including both instrument commands and other engineering and navigational activities. All of these commands are assembled and integrated well in advance and run through a rigorous series of constraint checks and simulation procedures to ensure proper execution and the continued health and safety of the spacecraft.

A sequence will typically span multiple spacecraft orbits of Saturn, also referred to as “revs.” Each rev by definition contains one apoapse and one periapse. Sequences are further broken down into individual “segments,” the responsibility for which falls to a given TWT or OST to “integrate.” Integration consists of the initial process of negotiation between instruments, and establishes the basic timeline and operational parameters of the segment. There are six types of segments corresponding to the different TWT/OSTs: Saturn, rings, MAPS (magnetosphere and plasma science), TOST (Titan), SOST (other satellites) and “XD,” or cross-disciplinary, which generally refers to periods requiring longer viewing distances from the desired target. Note that the segment type does not necessarily rule

out observations of other targets during that time period. Appendix C provides a breakdown of the Cassini mission chronology by orbit/rev and sequence.

Within a given segment are individual observations, typically designed by a single scientist for a specific purpose, and consisting of one or more “triggers” which in turn corresponds to a specific set of camera commands. These commands are contained in files – called Instrument Operations Interface files, or IOIs – the creation of which will be discussed in more detail later in Section 2. Note that there is a one-to-one relationship between IOIs and triggers: all IOIs have a corresponding trigger, and vice-versa. Triggers also have relative timing, which means they can be re-used over the course of the mission.

From an uplink perspective, it is useful to think about a given ISS observation as a specific combination of geometry, timing, and camera parameter settings. The geometry information, as well as timing to some extent, is specified through the use of a proprietary software tool called PDT, the Pointing Design Tool. PDT reads in the projected spacecraft and target ephemerides (contained in files called SPICE kernels, to be discussed below) and creates a simulation of the ISS view at a specified set of times. This allows the user to plan the necessary slew and framing events for the time period in question while taking into account the long list of flight rules and pointing constraints. PDT outputs a pair of files (a C-kernel and a Short-Form Output File, or SFOF) containing the pointing and timing associated with each spacecraft movement in an observation. This information is then read into another piece of software - usually the ISS Pre-Commanding Tool (ISSPT), discussed below – for the next step of calculating of exposure times, image compression, predicted data volume, and all of the various ISS camera settings, and writing these to an IOI file.

Figure 1 illustrates the three main coordinate systems used to describe the instrument and image orientation, and Figure 2 shows the relative instrument fields of view. Note that both the NAC and WAC are pointed along the spacecraft's -Y axis; they are boresight-aligned, with each other as well as

with Cassini's other optical remote-sensing (ORS) instruments, CIRS, UVIS and VIMS; and that the optics for both cameras produce a 180-degree rotation about the boresight at the focal plane.





Figure 1: Relationship between ISS spacecraft and image coordinate systems. From Porco et. al. (2004).

In terms of the values recorded in ISS VICAR image files (see Section 3), increasing sample $(+X_{im})$ direction corresponds to $+Z_{sc}$, and increasing line $(+Y_{im})$ to $-X_{sc}$. In order to recover the true orientation of the image as projected onto the sky, it is therefore necessary to flip the VICAR image array about the X_{im} -axis, which is to say, display the image such that the (line, sample) origin point is at top left.

The Cassini spacecraft orientation is controlled by reaction wheels, which provide exceptionally stable viewing for the ISS cameras. In star images taken during the cruise phase of the mission, exposures taken over the course of 50 minutes exhibited a pointing variation of only $18 \mu\text{rad}$, or about 3 NAC pixels. Relative pointing accuracy is $\sim 50 \mu\text{rad}$ in both the line and sample directions.

PROJECTION ON SKY (ALONG Y-AXIS)



*ISS NAC calibration shows offset by +0.822 mrad in X and -0.426 mrad from in Z from the -Y axis. Other instruments are shown in their nominal positions.




Figure 2: Fields of view of optical remote sensing instruments on Cassini. From Porco et. al. (2004).

The PDT designer has many options available for tracking targets during an observation. The ISS cameras can track on the target center, on a specific latitude and longitude point, on a particular point on the limb, or any variation on these. Tracking can be co-rotational or stationary with respect to rotation. Ultimately, the tracking rate is constrained by the AACS (Attitude and Articulation Control Subsystem) mode, which is set to either RWA (Reaction Wheel Assembly) or RCS (Reaction Control Subsystem). The RCS has a faster maximum tracking rate than the RWA, at the expense of slightly worse pointing accuracy, and is often employed during close flybys or any situation requiring

additional control over the spacecraft trajectory. The RWA is by far the more commonly-used AACCS mode used during ISS observations.

Beginning in sequence S37, mission scientists started noticing deterioration in the performance of the reaction wheels, and implemented more stringent pointing requirements designed to reduce wear on the RWA. As a general rule, spacecraft rotations in the orbital plane are unproblematic, whereas rotations outside of the orbital plane can cause the reaction wheels to pass through high-friction, low-RPM regions that degrade performance over time. As a result, the new flight rules limit observations that demand multiple small pointing corrections (such as AZSCANS; see Appendix B) and require long slews of over 60 degrees to execute 30% slower.

The inertial attitude of the spacecraft is determined during flight through analysis of star positions measured by the Stellar Reference Unit (SRU) or “star tracker” affixed on the Remote Sensing Palette. Up to five star positions are used in the navigation calculation, resulting in an absolute pointing accuracy of 60 μ rad in the SRU focal plane, or about 10 NAC pixels. However, because the SRU is oriented in the spacecraft +X direction, orthogonal to the ISS boresights, this pointing accuracy is only achieved in the ISS line direction, corresponding to rotation about the spacecraft Z-axis. In the ISS sample direction, corresponding to rotation about the spacecraft X-axis, pointing must be determined by measuring twist about the SRU boresight, the accuracy of which strongly depends on the distribution of stars in the SRU field of view. As a result, absolute pointing of the spacecraft in the sample direction is limited to about 900 μ rad, or 150 NAC pixels.

2.3 Camera Parameterization Part 1: Camera Settings and Operation

The purpose of this section is to provide a description of the various camera settings and options

that are available to the ISS, and some context for their use. While a brief summary of the camera hardware is included here, a far more detailed accounting of the Cassini ISS optics and electronics can be found in Porco et. al. (2004), and is recommended reading.

The ISS Narrow-Angle Camera (NAC) is an $f/10.5$ reflecting telescope with an image scale of $\sim 6 \mu\text{rad}/\text{pixel}$, a 0.35×0.35 degree field of view, and a spectral sensitivity from 200-1100nm, which is spanned by 24 filters, arranged in two filter wheels of 12 filters each. The Wide-Angle Camera (WAC) is an $f/3.5$ refractor with a $\sim 60 \mu\text{rad}/\text{pixel}$ image scale, and a 3.5×3.5 degree field of view. The WAC's refractive optics limit its UV sensitivity such that the overall spectral coverage is 380-1100nm. The WAC has 18 filters arranged in two filter wheels of 9 filters each. Geometric fidelity in both cameras is very good across the field of view, but better in the NAC, again due to its optical design and narrow viewing angle. Ground-based and in-flight measurements indicate geometric distortion of less than a pixel in the corners of the NAC field of view, and about 3.36 pixels in the WAC.

Filters

In both cameras, the filter wheels are arranged in-line for greater flexibility – allowing for a polarizer and a color filter, say, or the bandpass created by two overlapping color filters. Each filter wheel thus also has a “clear” filter for when only one (or no) filter is to be used. The complete list of filters and a brief scientific justification for their inclusion is provided in Table A.1 of the Appendix, and the basic bandpass parameters for the most commonly-used filter combinations (and some uncommonly-used combinations) are given in Table A.2.

Measurements of the point-spread function (PSF) are also provided in Table A.2. Point sources imaged by the NAC and WAC have 4 and 6 diffraction spikes, respectively, caused by their internal optics support structures. PSF width is less than 1.5 pixels for most filters, with one notable exception

being the WAC IR2, IR1 combination, with a full-width at half-maximum of nearly 5 pixels. This is explained by the WAC's refractive optics, which are unable to achieve perfect focus across the entire spectral sensitivity range. To compensate, the thicknesses of the WAC clear filters were chosen so as to optimize focus when combined with the filters in the opposite wheel; the IR2, IR1 combination cannot take advantage of this consideration. The other notably anomalous PSF is that of the NAC GRN filter, whose point-spread response contains a “ghost” peak a few tens of pixels from the main peak, with an amplitude approximately 1% that of the main peak.

As discussed in West et. al. (2010), in-flight images of bright stars and satellites have been used to characterize the PSF out to the edge of the detector for most filters. This extended PSF contains significant excess flux, several orders of magnitude above what would be expected of a theoretical diffraction-limited point source, most likely caused by internal reflections and off-axis stray light.

The ISS filter selection is determined by the science objectives, the target selection, the capabilities of the instrument, and the nature of the mission itself, which is characterized by highly eccentric Saturn-centered orbits and very fast flybys, and thus requires coverage of a wide range of viewing and phase angles, sometimes in very short time intervals. The fastest flybys occur during periods of closest approach, and in these geometries – often offering unique viewing opportunities that won't be duplicated during the mission - the NAC field of view is sometimes too narrow to provide for full spatial coverage of the desired target in the available time. To help remedy this, more than half the NAC filters have been duplicated in the WAC, to allow for broad spectral coverage in both narrow and wide-angle contexts. These include seven medium and broad-band filters for spectrophotometry, two methane and two continuum-band filters for atmospheric vertical sounding, two clear filters, and a narrow-band H α filter for lightning observations.

The clear filters allow transmission across the entire spectral range of the instrument, and thus provide the greatest sensitivity for imaging faint objects, or when short exposures are required due to a

tight observing window, or to minimize smear during close flybys. In other words, they are often used when maximizing signal-to-noise is the primary consideration.

The medium and broad-band filters allow for color imaging across a wide range of targets and imaging situations. They are BL1, GRN, RED, IR1, IR2, IR3, and IR4 (common to both cameras), UV1, UV2 and UV3 (NAC only), and VIO and IR5 (WAC only). Only the NAC has UV sensitivity due to its reflective optics, and a special Lumogen coating on the CCD, which was unavailable at the time of the Voyager and Galileo missions. The NAC takes advantage of this new capability by splitting up its UV coverage over three filters, which together provide visibility of targets such as stratospheric aerosols, auroral phenomena, and ring and satellite materials of special interest. However, solar irradiance at Saturn is extremely low in the UV, so very long exposure times are generally necessary, especially for UV1.

Methane-band and corresponding continuum-band filters (MT1 and CB1, MT2 and CB2, MT3 and CB3) are included primary for atmospheric and surface studies of Titan, allowing the ISS to probe more deeply through its layers of haze than would otherwise be possible. The CB3 filter in particular, with its effective wavelength well into the near-IR, provides the ISS with its best look at Titan's surface and tropospheric clouds.

Sets of polarized filters on both cameras allow for polarimetry observations across a range of targets and viewing situations. The NAC's first filter wheel carries three polarizers – P0, P60, P120 – which are sensitive to visible wavelengths, and whose principal transmission axes oriented at 60 degrees to one another. This arrangement is the most efficient one that allows for remote measurement of intensity, degree and direction of linear polarization from any camera orientation. While the full usable wavelength range of the visible light polarizers extends from around 350 nm to 750 nm, they are most effective from 450 nm to 650 nm, where the principal axis transmission is between 0.45 and 0.65, and orthogonal transmission is less than 1%.

The WAC, having fewer filter wheel slots than the NAC, carries only two polarized filters in a perpendicular orientation – IRP0 and IRP90 – which provide for measurement of intensity and the Stokes parameter, Q , at infrared wavelengths. The NAC also carries a single polarized filter sensitive to infrared wavelengths, IRP0, which can be used in conjunction with non-polarized clear filter images (CL1, CL2) to measure Q , though with less accuracy than is possible with two orthogonal polarizers. The IR polarizers exhibit much better performance over their range of spectral sensitivity (700 nm – 1000 nm) than the NAC visible polarizers, with a principal transmission greater than 0.9, and orthogonal transmission of 0.001 or less.

Shutter

Exposure of the ISS is controlled by the shutter assembly. This two-bladed system is located at the focal plane of each camera, and oriented such that shutter movement occurs in the image sample direction (that is, the blades are kept parallel to the CCD columns), and operates in a 3-step process: 1) both blades “reset” by moving across the detector to their home position in preparation for exposure; 2) the first blade sweeps linearly across the detector in the negative sample direction to “open” the shutter; 2) the second blade follows the first in the same direction to “close” the shutter at the end of the commanded exposure time. Shutter movement can also be inhibited completely during an exposure to allow for the measurement of dark current.

There are 63 discreet exposure time settings available, ranging from 0 to 1200 seconds, and one “No-op” (no operation) setting in which no shutter movement or CCD readout occurs. The smallest non-zero commanded exposure time is 5 ms. Because the shutter is mechanically activated, actual shutter times differ from commanded times in both cameras. This so-called shutter offset has a constant component as well as a sample-dependent component. The constant component is about -2.85 ms in the NAC (which is to say, a commanded exposure time of 5 ms yields an actual exposure of 2.15 ms), and

-2.86 ms in the WAC, with an uncertainty of ± 0.25 ms. See Table A.3 in the Appendix for the full set of commandable exposure times.


Detector

Both cameras feature identical 1024 x 1024-pixel CCD detectors with pixels 12 μm to a side, fronted by a fused silica window. Each CCD is capable of on-board pixel summation (2x2 or 4x4, in addition to the standard 1x1), which can be useful for either increasing signal-to-noise or decreasing data volume. There are four gain modes (shown in Table A.4), three of which were designed to correspond to the full-well capacities of the different summation modes, and one of which is an additional high-gain mode intended for faint targets.

The response of the CCD detectors to incident light is determined by their respective quantum efficiencies as a function of wavelength, i.e. the number of electrons released in a given pixel's potential well per incident photon. The quantum efficiency in both cameras plateaus at ~ 0.35 e⁻/photon between 550 and 750 nm. Spectral response falls off steeply towards the infrared, and disappears completely at 1100 nm. Moving towards the blue end, sensitivity drops to about 0.14 between 550 and 475 nm, and remains flat down to 200 nm, the benefit of a Lumogen CCD coating. Note that while both detectors are nominally sensitive to these short wavelengths, the WAC optics are opaque to UV light.

During readout, the entire bottom line of the CCD is shifted downward into the serial register and then is read out in the sample direction, pixel by pixel, into the signal chain, followed by the next line, etc. This analog signal passes through the CCD on-chip amplifier and is then encoded to a 12-bit digital quantity – known as a data number or DN – by the analog-to-digital converter. (Although these DNs contain 12 bits of information, they are stored on the spacecraft as 16-bit values; the upper four bits are ignored.) The camera operator then has a choice of data conversion: either keep the original 12 bits (no conversion) resulting in DN values from 0 to 4095, or convert to 8 bits, resulting in DN values

from 0 to 255. There are two types of 8-bit conversion: the first simply truncates each pixel DN to its least-significant 8 bits (alternatively known as LS8B or 8LSB) allowing all values higher than 255 to “wrap” back around to 0; and the second converts to 8 bits by use of a look-up table (LUT) that applies a variation of square-root encoding (see Figure 3). This type of conversion matches the quantization level to the photon noise, thus spreading information content more evenly across the DN range, and is generally the preferred method for conserving data volume when full dynamic resolution is unnecessary. This quantization does however make it more difficult to remove sources of low-level background noise like the 2 Hz noise (see Section 4). LS8B mode is most useful when imaging faint targets where DN values are not expected to go above 255, or for imaging smoothly-varying targets that can later be “unwrapped” in post-processing. Complexly-varying targets with large DN values that have been “wrapped” by LS8B mode are very difficult if not impossible to recover.



β

Figure 3: Illustration showing the modified square-root function used to convert 12-bit data to 8 bits in LUT mode. From Porco et. al. (2004).

The CCD full-well capacity is approximately 120,000 e⁻/pixel, and the summation well capacity is ~1.6x10⁶ e⁻/pixel. Note that the summation well is a separate entity and its capacity does not simply

scale with the number of summed pixels. For this reason, ground calibration analysis found significant non-linearity in 4x4 mode with gain state 0 for DN values greater than about 1000. For all other gain/summation states, deviation from linearity was found to be less than 1% for data number (DN) values greater than about 500 in both detectors.

It's important to understand the difference between well saturation and DN saturation, both of which will destroy the photometric utility of a pixel measurement. DN saturation is simply when the pixel DN level reaches the limit established by its bit encoding: 255 for 8-bit data, and 4095 for 12-bit. Full-well saturation occurs when a pixel reaches its well capacity, allowing excess electrons to spill over into adjacent pixels. This spill-over, or “blooming,” tends to occur along the column direction, and in cases of extreme overexposure can result in a significant or total loss of data from those columns. To help minimize this problem, both detectors are equipped with an “anti-blooming” feature that reduces this effect. At the beginning of the mission, anti-blooming mode was set to “ON” for most ISS observations by default, but concerns about side-effects – notably bright/dark pixel pair artifacts, and other changes in the detector noise behavior – eventually led scientists to restrict its use to cases where pixel saturation is expected.

Compression

The amount of data ISS can acquire in a given observation is determined by the length of the observing window and the available data volume. The latter is constrained both by the size of the spacecraft's solid-state recorder (SSR) and the limited communication bandwidth available for returning data to Earth. It is often therefore necessary to conserve data volume as much as possible. This is accomplished primarily via three camera settings, two of which – 8-bit conversion and summation mode – have already been discussed. The third is compression. The ISS has both lossy and lossless compression options, as well as the option to choose neither (although this is seldom done in

practice). Both compression algorithms are implemented in hardware and are documented in detail in Porco et. al. (2004) and elsewhere.

To summarize, lossy compression is a variant on the familiar Joint Photographic Experts Group (JPEG) algorithm, and can achieve very high compression ratios, mostly by removing information at high spatial frequencies. Although the compression settings are configurable to minimize information loss for different kinds of targets, lossy compression is generally avoided when taking images for photometry. Beginning with the start of the CSM/XXM, use of lossy compression mode has been discontinued in order to save on work-hours required for post-processing and image reconstruction.




Figure 4: Detail of a 4x4-summed image of Phoebe taken during approach showing line truncation caused by the lossless compression algorithm.

Lossless compression is based on Huffman encoding, a high-efficiency numerical scheme in which the length of the bit sequence used to encode a number is based on the frequency of that number's occurrence. As its name implies, losslessly-encoded images can be uncompressed and recovered on the ground with no information loss whatsoever, provided the image entropy (a measure of the amount of visual information in a scene) in a pair of lines does not exceed the level where 2:1 compression is achieved. That is to say, 2:1 is the minimum compression ratio when in lossless mode.

In images where the amount of visual information is too high for it to be losslessly compressed to half its original size, the algorithm will truncate as many pixels as it needs from the end of each pair of lines to meet this requirement. See Figure 4 for an example.

Image Event

Table A.5 in the Appendix gives the steps involved in a single image event, and their order of operations. While an image event usually produces a single image by a single camera, the instrument can also be commanded to execute a simultaneous exposure by use of the “BOTSIM” (both simultaneous) setting. The total time in which an image event occurs is called a framing time, and it is composed of a prepare cycle and a readout cycle. The term “prepare cycle” is a bit of a misnomer, because it actually contains every step in the imaging process (including the exposure itself), with the exception of the image readout.

An imaging event begins with shutter reset from the previous exposure. This is followed by movement of the filter wheels to their commanded positions – first the NAC, then the WAC – and then the light flood and erase, which involves flooding the CCD with light to ~50 times the saturation level, followed immediately by a readout. This step is meant to eliminate variation in the residual bulk image (RBI, to be revisited later in the context of dark subtraction), and helps to ensure a uniform, repeatable starting condition for the detector. Next comes the exposure itself; note that in BOTSIM mode, the two exposures are synchronized at the closing of the shutter. When exposure is complete, the CCDs are read out in sequence, NAC first, then WAC.

The total times allotted to the prepare and readout windows are, like the exposure times, only available in discreet increments, of which there are 16 for each. Tables A.6 and A.7 in the Appendix list these prepare and readout indices and their corresponding times. The minimum possible prepare window is 4.475 seconds (for a single-camera exposure of less than 2 seconds with the filter wheel

moving no more than 3 positions) and the maximum is 1212.475 seconds (to accommodate a 1200-second BOTSIM exposure with the largest possible filter wheel movement).

The readout window behaves a little differently, as there are only four possible readout times per camera, but these scale with the spacecraft data telemetry rate. The possible telemetry rate settings and values are given in Table A.8. Assuming a telemetry rate of 24 packets per second, the most common rate used during ISS observations, the smallest readout window is 12.525 seconds per camera and the largest is 100.525 seconds. Combining the prepare and readout windows to get a full framing time, we see that the minimum possible amount of time between ISS imaging events at the standard telemetry rate is 17 seconds, which is enough time to read out a standard 12-bit lossless 2x2-summed image. (This can be reduced to 11 seconds if the highest telemetry rate of 48 packets/sec is used.) By contrast, a typical 12-bit lossless 1x1 image – commonly used for spectrophotometric and calibration observations – will require a readout window of 50.525 seconds, and a total framing time of 55 seconds at the standard telemetry rate.

While the prepare window is completely determinate, the readout window is not, and must be chosen carefully to accommodate the full readout time required for the image. This requires predicting the compression ratio in advance. Because this is not always possible to great accuracy, the camera commanding tool ISSPT includes significant buffer in its compression predictions.

Compression settings are not the only camera parameters that affect readout rate. In fact, most camera settings affect readout behavior to some extent. Because of this, estimating the line-dependent readout time for a given set of camera parameters – crucial for both observation planning and calibration – is a rather non-trivial task requiring a software algorithm designed for this purpose. This algorithm, called “linetime,” was written by Vance Haemmerle at JPL and is included in both the CISSCAL calibration software and ISSPT.

2.4 Camera parameterization Part 2: The ISS Pre-commanding Tool (ISSPT)

It is useful for anyone learning about ISS data to have a basic understanding of the ISS Pre-commanding Tool, which is the primary software tool used by Cassini Imaging scientists to design camera commands. Without getting into too much specific detail, this section will lay out the basic structure and purpose of the software to give a sense of how it is used in practice. The considerations that have gone into the design of ISSPT are crucial for understanding use of the cameras in general, and the data that results.

The purpose of ISSPT is to facilitate the selection of camera parameters for a given camera pointing in a given period of time. As discussed above in Section 2.2, its input is the camera pointing geometry (contained in a C kernel file) and timing information (contained in an SFOF) produced by PDT for a particular observation; its output is the Instrument Operations Interface (IOI) file, containing the complete list of camera parameters (40 in all, not including global parameters), arranged in “KEYWORD=VALUE” format, corresponding to every imaging event in the observation (or, more accurately, in a subset of the observation called a “tracking period,” during which a specific camera trigger is to be executed).

Once ISSPT has the timing information for each pointing in the observation (and each slew in between pointings), it uses the camera orientation, spacecraft trajectory and solar system body ephemerides to reconstruct, via ray tracing, a “footprint.” Footprints are essentially simplistic predictions of the image scene, though they contain a wealth of information, including target distance, incidence and emission angle, solar phase angle, and a thumbnail image from which can be calculated the relative fraction of each type of surface or object contained in the field of view (e.g. 20% Saturn atmosphere, 40% rings, 40% dark sky, each with associated lossy compression settings).

Having simulated a footprint at each pointing, ISSPT is then ready to apply a target model

relating geometry to target brightness. The target model is chosen by the user and determines the incident flux level on which the exposure selection will be based. There are separate models for Jupiter, Saturn, Titan, the icy satellites, the main A, B and C rings, and the faint D, F and G rings.

Alternatively, the user can bypass the target model calculation and simply input an I/F value or spectrum. For stellar observations, there is a star target model that accepts stellar type and magnitude.

At this point ISSPT has all of the information it needs to calculate exposure times and prepare windows for each exposure in a footprint. The exposure calculation is designed to be conservative, though the user may set the target DN level and full well level to be higher or lower as desired. Once a nominal exposure time has been chosen, ISSPT finds the resultant expected DN level and uses it, along with information about the footprint scene, to estimate the lossy or lossless compression ratio, and therefore data volume, for each image. Finally, readout time is computed, and an appropriate readout window selected. After that, it is up to the user to construct the specifics of the observation by adding or subtracting exposures and footprints, altering camera parameters from their defaults, and adjusting timing.

Timing adjustment is accomplished in ISSPT by use of a timing table which lists all the image events in each footprint and their timing relative to the trigger start time, and allows the user to add wait times in between images, check for warning or error conditions such as readout or exposure during slew, and take advantage of a built-in “Loop” feature in the cameras that simplifies commanding in situations when you want to repeat an exposure or set of exposures many times within a pointing. In the simplest cases, a single camera or BOTSIM exposure may also be repeated up to 255 times by setting the “iteration count” parameter to greater than 1.

Finally, ISSPT allows the user to set image metadata, much of which will eventually end up in the image labels and other data products like the archive index.tab file (discussed in the Section 3).

ISSPT automatically records its software version number, the selected target model, and the name of

the observation to the METHOD_DESCRIPTION keyword, and allows the user to include additional notes as desired. There are also keywords specifying the image type - science, navigation, calibration, engineering – and for support observations, an additional flag indicating which instrument team the images are intended for. Other non-camera parameter values that are calculated by ISSPT and recorded in the metadata include predicted compression (in bits/pixel, included in the INST_CMPRS_RATE keyword), predicted data volume (EXPECTED_PACKETS), and predicted percentages of full-well capacity and maximum DN (EXPECTED_MAXIMUM).

The observation planning process is an iterative one, and a given observation may pass through PDT and ISSPT several times before the final pointing and camera design are settled upon. Multiple constraint checks at different points in the process check for errors in timing and data volume calculations. When the design is complete and has passed all checks, the camera commands from all the IOIs/triggers to be executed in the sequence are merged with all of the spacecraft pointing commands, and everything is finally converted to Instrument Expanded Blocks (IEBs), the data format that is actually uploaded to the spacecraft. There is one IEB load uplink per instrument per sequence.

The following section will provide a detailed examination of the fruits of these efforts: the ISS images themselves, and a summary of all imaging conducted over the course of the Cassini mission thus far.

3. The ISS Data Set

3.1 ISS Science Data: Summary and Search Methods

The ISS image data consists of both the image itself - brightness as a function of line and sample, in units of raw data numbers (DNs) - and an associated label, which contains image-specific metadata in the form of “KEYWORD=VALUE” pairs. A thorough description of the relevant file formats and their contents is given in the subsections that follow, but first we will concentrate on a more generalized view of the ISS data as a whole to aid the first-time user in developing a conceptual basis for searching. In this discussion, terms in all caps with underscores in place of spaces (“ALL_CAPS”) refer to specific PDS keywords that describe the images, and are available for searching, either in the VICAR image label, the detached PDS label, or the index.tab files included with the PDS data archive volumes. Of these, the index.tab contains the most complete set of searchable keywords, including geometry keywords derived from the images themselves.

There are a number of ways to subdivide and organize the ISS data set that are conceptually useful. Perhaps the most basic subdivision is simply which of the two cameras is used – an INSTRUMENT_ID of “ISSNA” refers to images taken with the NAC, and “ISSWA” to images taken with the WAC. (The camera is also indicated by the first letter of the image filename, as described below.) As of this writing in the summer of 2011, the Cassini ISS cameras have together acquired over 280,000 images, around two-thirds of which are NAC images.

Each image has an IMAGE_OBSERVATION_TYPE which can be set to one or more of the following values: “SCIENCE,” “CALIBRATION,” “ENGINEERING,” “OPNAV” (operations/navigation), and “SUPPORT IMAGING.” The vast majority of images of use to scientists will be of type “SCIENCE” or “SUPPORT IMAGING”; images taken for “OPNAV” or

“ENGINEERING” are typically only used for in-flight support of the spacecraft during the mission, and “CALIBRATION” images are used for assorted in-flight image calibration tasks, to be discussed in Section 4.

A “support image” is an image requested by an instrument team other than ISS. As shown in Figure 2, ISS shares a boresight ($-Y_{sc}$) with the other Optical Remote Sensing (ORS) instruments UVIS, CIRS, and VIMS, each of which occasionally requests simultaneous ISS support imaging for navigation and context.

As discussed in Section 2, the Cassini mission is split up into “sequences” of commands, which are compiled together and periodically uplinked to the spacecraft. A sequence typically contains two to three orbits of Saturn, or “revs,” and each rev is typically split up into a number of different “segments,” chunks of time that are divvied up amongst the different TWTs and OSTs during the science planning process. Within each segment, individual instruments are allotted time for their observations, and data volume according to the number of downlink periods within that segment. Each observation is given a unique identifier, the observation request name, which is recorded in the image label as OBSERVATION_ID and has the format:

$$[INST]_[REV][TI]_[UNIQUENAME]_[PINST]$$

where *INST* is the abbreviation for the instrument being commanded (always “ISS,” in our case), *REV* is the rev number, *TI* is a 2-character target identifier (see Table 3), *UNIQUENAME* is a unique descriptor chosen by the observation creator, and *PINST* indicates the name of the prime instrument if other than the ISS; if the ISS is prime, *PINST* is “PRIME” (in the nominal mission and CEM) or “PIE” (in the CSM). So for example, an OBSERVATION_ID of “ISS_039RI_PHOTLIT001_PRIME” indicates an ISS-prime rings observation in rev 39, “ISS_139RH_REGMAP001_PIE” is an ISS-prime Rhea observation in rev 139, and “ISS_192TE_LOPHASE001_UVIS” indicates a UVIS-prime Tethys observation in rev 192. Observations in which the commanded instrument is not prime are called

“ride-alongs” because the commanded instrument is “riding along” with the prime instrument. ISS ride-along observations may or may not also contain support images.

The *UNIQUENAME* portion of the *OBSERVATION_ID* is left up to the observation creator and thus follows no fixed convention, but it can often be quite descriptive, particularly for certain types of observations that are repeated many times over the course of the mission. In the ISS-prime examples above, one can surmise that “RI_PHOTLIT” corresponds to photometry of the lit side of the rings, and “RH_REGMAP” refers to regional mapping of Rhea. Appendix B of this document provides a partial list of observation types indicated by the *UNIQUENAME* portion of the *OBSERVATION_ID*, and organized by target discipline.

AG	Aegeon	FT	Flux tube	PH	Phoebe	RI	Rings (general)
AN	Anthe	GA	Ganymede	PL	Pallene	SA	Saturn
AT	Atlas	HE	Helene	PM	Prometheus	SC	Spacecraft Activity
CA	Callisto	HY	Hyperion	PN	Pan	SK	Skeleton Request (for mission planning)
CO	Co-rotation	IA	Iapetus	PO	Polydeuces	SR	Spacecraft RAM direction
CP	Calypso	IC	Instrument Calibration	PR	Plasma RAM direction	ST	Star
DA	Daphnis	IO	Io	RA	A-ring	SU	Sun
DI	Dione	JA	Janus	RB	B-ring	SW	Solar Wind
DN	Downstream of the wake	JU	Jupiter	RC	C-ring	TE	Tethys
DR	Dust RAM direction	ME	Methone	RD	D-ring	TI	Titan
EA	Earth	MI	Mimas	RE	E-ring	TL	Telesto
EN	Enceladus	NA	Not Applicable (used for planning)	RF	F-ring	TO	Io torus
EP	Epimetheus	OT	Other	RG	G-ring	UP	Upstream of the wake
EU	Europa	PA	Pandora	RH	Rhea		

Table 3: Complete list of target identifiers in OBSERVATION_ID.

In terms of their pointing strategies, observations tend to fall into one of a few different categories. A “point-and-stare” involves pointing at a target or a particular location on a target and tracking that position over some period of time. The simplest example of this type of observation

would consist of a single footprint. Multiple evenly-spaced exposures of a fixed target taken over a relatively short period of time is usually referred to as a “movie,” and may have “MOV” or “MOVIE” in the OBSERVATION_ID. An observation that consists of multiple overlapping footprints offset from each other is called a “mosaic” (see Figure 5). A mosaic can be as simple as a two overlapping footprints when one footprint isn't enough to cover the entire target, or as complicated as a complete azimuthal scan of the rings at a fixed ring radius, requiring hundreds of individual pointings with fine offsets between each overlapping footprint. (See “AZSCAN” in Appendix B and Figure 6.) There are also movies composed of individual mosaics, and within a given observation, just about any other combination of movies, mosaics, scans and stares, depending on the target geometry, the available data volume and time window, and the ultimate science goals.




Figure 5: Example of a typical NAC mosaic of Saturn visualized by the Pointing Design Tool (PDT). This particular observation is a 2x2 movie in multiple filters, intended for the study of atmospheric dynamics. From Porco et. al. (2004).




Figure 6: Example of a NAC azimuthal scan (AZSCAN) observation of the outer B-ring edge and Cassini Division at moderate phase and high sub-spacecraft latitude.

Here, UTC refers to the widely-used “Coordinated Universal Time” standard, and conforms to the CCSDS ASCII time code format:

yyyy-dddThh:mm:ss.fffZ

(Example: “2011-009T12:00:00.000Z”; here, ddd corresponds to the day of year.) With the exception of IMAGE_NUMBER and PRODUCT_ID, all of these time keywords reflect the full spacecraft clock resolution of one subRTI, or 1/256th of a second (approximately 4 msec). When the shutter is inhibited (i.e. SHUTTER_STATE_ID = “DISABLED”), the IMAGE_MID_TIME = START_TIME = STOP_TIME, and all three represent the start of the exposure window during the prepare cycle of the image. In the case of simultaneous NAC and WAC imaging (BOTSIM) the time of shutter close will be

identical for both cameras, and thus also the IMAGE_TIME, IMAGE_NUMBER, SPACECRAFT_CLOCK_STOP_COUNT, and STOP_TIME. The PRODUCT_ID, however, will remain unique due to the “N” or “W” prefix.

Keyword	Description
SPACECRAFT_CLOCK_STOP_COUNT	The seconds portion of the spacecraft clock at shutter close
SPACECRAFT_CLOCK_START_COUNT	SPACECRAFT_CLOCK_STOP_COUNT – EXPOSURE_DURATION
IMAGE_TIME	SPACECRAFT_CLOCK_STOP_COUNT, converted to UTC format
IMAGE_MID_TIME	SPACECRAFT_CLOCK_STOP_COUNT – (EXPOSURE_DURATION/2), converted to UTC
IMAGE_NUMBER	The seconds portion of the spacecraft clock at shutter close
STOP_TIME	SPACECRAFT_CLOCK_STOP_COUNT, converted to UTC
START_TIME	SPACECRAFT_CLOCK_STOP_COUNT – EXPOSURE_DURATION, converted to UTC
PRODUCT_ID	Unique identifier for each image, constructed as follows: SPACECRAFT_CLOCK_CNT_PARTITION + “_” + [I]SPACECRAFT_CLOCK_STOP_COUNT, where I is “N” for the NAC, and “W” for the WAC
PRODUCT_CREATION_TIME	Time corresponding to creation of image file on the ground, in UTC

Table 4: Image keywords related to timing.

For searches based on target and target geometry, there is a wide assortment of keywords that may be used, most of which are found not in the image label but rather in the index.tab file located in the index/ directory of each data archive volume. Most of these target geometry keyword values are calculated after downlink, at the time of ingestion into the ISS image database. Further refinements to the pointing are made in a process known as “C-smithing” where ISS images are examined and “navigated” in order to reconstruct the pointing as accurately as possible after the fact. The software used in this process is called “Autonav,” and the end result of the navigation process is a set of updated target geometry keywords, as well as new C-kernels. From the ISS archsis.pdf file: “Autonav uses an array of object detection algorithms in conjunction with the most recent spacecraft position and

orientation kernels to navigate the images. The output of Autonav for any particular navigated image is a single, discrete C-kernel for the IMAGE_MID_TIME. [...] Though the success rate of Autonav is high, it is not 100% successful. The code was structured to minimize the number of false-positive navigations[...] However, all of these thresholds and verification steps do not absolutely prevent Autonav from producing false results, so future users are warned to exercise caution with respect to these [data]. Autonav results, when accurate, will greatly improve the accuracy of the geometric quantities calculated for the index.tab file.”

The C-smithed pointing information generated by Autonav is packaged up into C-kernel files covering approximately 30-day intervals and subsequently delivered to the PDS mission archive and made available in the “ck” subdirectory of the NAIF kernel repository. These files have names like “11152_11181ca_ISS.bc,” where the prefix numbers indicate the relevant time range in YYDOY format. Note that interpolation between pointings at discrete times is not implied – the updated pointing information is valid only for the IMAGE_MID_TIMES of the successfully navigated images within that time range.

If the computed Autonav accuracy is less than 5 pixels, the target geometry keyword values are populated with the values derived from the AACS pointing. (To determine whether the geometry keywords for a particular image were derived from Autonav or AACS pointing, one can look at the SPICE_PRODUCT_ID keyword giving the names of the SPICE kernels used in the geometry calculations; use of a C-smithed C-kernel created by Autonav will be indicated by the presence of a filename of the form <image filename> + '.bc', e.g. “W1691944231_1.IMG.bc”.) Targets which tend to prove problematic for Autonav include anything with a “fuzzy” outline, notably Titan and the faint rings. As the mission progresses, ISS/CICLOPS continues to work to improve the Autonav success rate, though as of this writing, improvements apply only to new images going forward and not ones that have already been archived with PDS. Any significant changes to Autonav are documented in the

errata.txt file.

There are a few different ways in which a user might want to search for the ISS data set for a particular target. The TARGET_NAME is the most straightforward keyword to use for this, as it simply identifies the primary target body at which the spacecraft was pointed during the observation. There is also the TARGET_DESC, which contains the name of the ISSPT target *model* on which the image exposure was based. Finally, and most comprehensively, is the TARGET_LIST generated by Autonav. This keyword contains a complete list of all target bodies falling within the field of view. Occultation is taken into account, such that a target is included if it has any part of its limb is not hidden by another body (or the rings). Queries based on TARGET_LIST should be used if the user is interested in identifying the broadest possible set of images featuring the desired target.

For more refined target searches, other geometry keywords populated by Autonav include RIGHT_ASCENSION and DECLINATION, the standard planetary viewing angles – INCIDENCE_ANGLE, EMISSION_ANGLE and PHASE_ANGLE – as well as target- and aimpoint-specific quantities like SUB_SPACECRAFT_LATITUDE and LONGITUDE, TARGET_DISTANCE and PIXEL_SCALE, and many others. A complete list of the target geometry-related keywords produced by Autonav is provided in Table 5.

Figure 7 gives an overview of Cassini's viewing geometry as a function of time for the entire tour of Saturn. Note periods of high sub-spacecraft latitude vs. periods where the spacecraft remains in the equatorial plane. The latter generally contain the majority of satellite flybys whereas the former contain more ring imaging, and high-latitude viewing geometries that are not possible from within the ring plane. A more detailed description of the mission geometry, along with a useful chronological summary of satellite encounters, can be found in the mission.cat file in the catalog/ directory of the archive volume.

Keyword	Description
CENTER_LATITUDE	Planetocentric latitude at aimpoint on target (or Ring_Aimpoint_Latitude if target is a ring) [deg].
CENTER_LONGITUDE	West longitude at aimpoint on target (or Ring_Aimpoint_Longitude if target is a ring) [deg].
CENTRAL_BODY_DISTANCE	Distance from spacecraft to center of Saturn [km].
COORDINATE_SYSTEM_NAME	The full name of the coordinate system to which the state vectors are referenced.
DECLINATION	Declination of camera optic axis [deg].
EMISSION_ANGLE	Emission angle at aimpoint on target [deg].
INCIDENCE_ANGLE	Incidence angle at aimpoint on target [deg].
LOWER_LEFT_LATITUDE	Planetocentric latitude of lower-left pixel [deg].
LOWER_LEFT_LONGITUDE	West longitude of lower-left pixel [deg].
LOWER_RIGHT_LATITUDE	Planetocentric latitude of lower-right pixel [deg].
LOWER_RIGHT_LONGITUDE	West longitude of lower-right pixel [deg].
MAXIMUM_RING_RADIUS	Maximum ringplane radius in image [km].
MINIMUM_RING_RADIUS	Minimum ringplane radius in image [km].
NORTH_AZIMUTH_CLOCK_ANGLE	Direction of the northward-pointing azimuth at the aimpoint on the target body [deg].
PHASE_ANGLE	Phase angle at subspacecraft point on target [deg].
PIXEL_SCALE	Size of one pixel at sub-spacecraft point on target body [km/pixel].
PLANET_CENTER	2-valued array: first, line number of target body center; second, sample number of target body center.
RIGHT_ASCENSION	Right ascension of camera optic axis [deg].
RING_CENTER_LATITUDE	Planetocentric latitude at aimpoint on ring [deg].
RING_CENTER_LONGITUDE	West longitude at aimpoint on ring [deg].
RING_EMISSION_ANGLE	Emission angle relative to target ring at aimpoint [deg].
RING_INCIDENCE_ANGLE	Incidence angle relative to target ring at aimpoint [deg].
RINGS_FLAG	If the target is a ring, then this quantity is "YES" if any part of that ring is visible in the image; "NO" otherwise. If the target is not a ring, then this quantity tests whether any part of the A, B, C, or D are visible in the image.
SC_PLANET_POSITION_VECTOR	3-valued array. X, Y, Z components of the position vector from the spacecraft to primary planet center, corrected for light-travel time and stellar aberration [km].
SC_PLANET_VELOCITY_VECTOR	3-valued array. X, Y, Z components of the velocity vector of primary planet relative to spacecraft, corrected for light-travel time [km/s].
SC_SUN_POSITION_VECTOR	3-valued array. X, Y, Z components of the position vector from the spacecraft to sun center, corrected for light-travel time and stellar aberration [km].
SC_SUN_VELOCITY_VECTOR	3-valued array. X, Y, Z components of the velocity vector of sun relative to spacecraft, corrected for light-travel time [km/s].
SC_TARGET_POSITION_VECTOR	3-valued array. X, Y, Z components of the position vector from spacecraft to target center, corrected for light-travel time and stellar aberration [km].

Keyword	Description
SC_TARGET_VELOCITY_VECTOR	3-valued array. X, Y, Z components of the velocity vector of the target relative to spacecraft, corrected for light-travel time [km/s].
SPICE_PRODUCT_ID	The names of the SPICE files used in processing the data.
SUB_SOLAR_LATITUDE	Planetocentric latitude of subsolar point on target; Ring_SubSolar_Latitude if target is a ring [deg].
SUB_SOLAR_LONGITUDE	West longitude of subsolar point on target; Ring_SubSolar_Longitude if target is a ring [deg].
SUB_SPACECRAFT_LATITUDE	Planetocentric latitude of subspacescraft point on target; Ring_Subspacecraft_Latitude if target is a ring [deg].
SUB_SPACECRAFT_LONGITUDE	West longitude of subspacescraft point on target; Ring_Subspacecraft_Longitude if target is a ring [deg].
TARGET_DISTANCE	Distance from the spacecraft to the center of the target [km].
TARGET_EASTERNMOST_LONGITUDE	Easternmost longitude visible on target body [deg].
TARGET_LIST	Name of each body visible in the image. A body is included if any part of its limb is not hidden by another body. Bodies may be obscured by rings. (Possible values: MERCURY, VENUS, EARTH, MARS, JUPITER, SATURN, URANUS, NEPTUNE, PLUTO, SUN, MOON, EARTH, MIMAS, ENCELADUS, TETHYS, DIONE, RHEA, TITAN, HYPERION, IAPETUS, PHOEBE, JANUS, EPIMETHEUS, HELENE, TELESTO, CALYPSO, ATLAS, PROMETHEUS, PANDORA, PAN, IO, EUROPA, GANYMEDE, CALLISTO, AMALTHEA, HIMALIA, ELARA, PASIPHAE, SINOPE, LYSITHEA, CARME, ANANKE, LEDA, THEBE, ADRASTEIA, METIS, SKY, MASURSKY, FOMALHAUT, SPICA, DARK SKY, NULL)
TARGET_NORTHERNMOST_LATITUDE	Northernmost latitude visible on target body [deg].
TARGET_SOUTHERNMOST_LATITUDE	Southernmost latitude visible on target body [deg].
TARGET_WESTERNMOST_LONGITUDE	Westernmost longitude visible on target body [deg].
TWIST_ANGLE	Twist angle of optic axis [deg].
UPPER_LEFT_LATITUDE	Planetocentric latitude of upper-left pixel [deg].
UPPER_LEFT_LONGITUDE	West longitude of upper-left pixel [deg].
UPPER_RIGHT_LATITUDE	Planetocentric latitude of upper-right pixel [deg].
UPPER_RIGHT_LONGITUDE	West longitude of upper-right pixel [deg].

Table 5: Target geometry keywords generated by Autonav software.

The Planetary Society also provides a useful timeline of the Cassini mission, including all targeted and non-targeted satellite flybys, at the following URL:

http://www.planetary.org/explore/topics/cassini_huygens/tour.html




Figure 7: Cassini mission geometry timeline.

The tables and links at the bottom of this page can be used to identify time periods for searching or to find the orbit number corresponding to a particular satellite encounter. (Recall that if a target and orbit/rev number are known, one can return all prime observations taken for that target by simply searching OBSERVATION_ID for values that match “ISS_[REV][TI],” where [REV] is the orbit number and [TI] is the two-letter target identifier from Table 3.) The times provided in these tables correspond to closest approach, so if one desires all ISS images from a particular flyby, it is best to define the search window to include at least 24 hours both leading up to and following the encounter.

Tables in Appendix C of this document give UTC times corresponding to each uplink sequence (searchable using the SEQUENCE_ID keyword) and rev/orbit in the mission.

3.2 Introduction to the PDS ISS Data Archive

Once an image is transmitted back to the Deep Space Network (DSN) and sent on to the Telemetry Data System (TDS) at JPL, it is reformatted by MIPL from a series of data packets back into a two-dimensional image. Recall that each 12-bit ISS image pixel is actually stored as a 16-bit value; in the reformatting process, the upper four bits in this unconverted, uncompressed data are converted from “1”s to “0”s. In addition, images that had been compressed, either losslessly or lossily, are automatically decompressed in the reconstitution process. All images are then sent to ISS/CICLOPS where they are ingested into the Archive Database, from which the ISS data archive is built.

Preliminary versions of images are generated immediately and distributed for instrument performance analysis. If there are missing data in the preliminary version, IO performs reconciliation in an attempt to create a more complete product. Once reconciliation is performed (within two weeks from time of downlink), a final version of the image is produced. Only the final versions of images are archived, and these are put on the archive volumes exactly as they are received from MIPL. The only additional processing that is done is to auto-navigate the images and physically assemble the archive volumes. Automated software is used to generate the archive disks by selecting the appropriate range of images, gathering the static documentation, and generating the index.tab file from the auto-navigation results. All further steps for processing, cleaning, and converting the images to physical units are left to the calibration pipeline, discussed in the following section.

The ISS archive is comprised of two types of archive volumes. The first is referred to as the “DATA” archive and is considered to be mostly static. These volumes contain the raw (uncalibrated) ISS experiment data record image files (EDRs) with their attached labels, detached PDS label files, helpful and required PDS files, and other useful documentation related to the image datasets. “Static” refers to the fact that, once produced and validated, the contents of these volumes are less likely to be

updated or modified. Exceptions are those cases where new information is made available or where images are reprocessed, due to previous errors, and are made available on later volumes. No calibration files are found on the DATA volumes, except in-flight calibration images, sequenced along with all other ISS images in spacecraft clock (SCLK) order.

The ISS team provides the imaging EDRs in raw, uncalibrated form only. In order to allow future users of these data to perform their own calibration processing, ISS has also made available the “CALIBRATION” archive. These archive volumes contain all calibration-related files, except in-flight images as mentioned above, along with support documentation, algorithms and software. The first ten volumes of the calibration archive contain ground-based images taken during pre-launch instrument testing, as well as the ground calibration report (contained in the document/report/ subdirectory) detailing analysis of these images and the derivation of various parameters important for calibration. Many of the algorithms and data files used in calibrating ISS images have been derived from this original ground calibration analysis.

The eleventh calibration volume (VOLUME_ID = “COISS_0011”) is dynamic, with ISS releasing periodic updates to the PDS. The calibration data files and CISSCAL software algorithms on this volume will continue to evolve and improve as knowledge of the instrument improves.

The ISS archive collection is further organized into three datasets:

1. The raw EDR images, in spacecraft clock (SCLK) order, from launch up to the start of the Saturn approach science phase. In addition to the EDR images, also included are the support images, in-flight calibration images and images used for navigation purposes.

Products in this dataset are identified with the following DATA_SET_NAME,

DATA_SET_ID and STANDARD_DATA_PRODUCT_ID:

CASSINI ORBITER EARTH/VENUS/JUPITER ISSNA/ISSWA 2 EDR V1.0
CO-E/V/J-ISSNA/ISSWA-2-EDR-V1.0
ISS_E/V/JEDR

2. The raw EDR images (in SCLK order), support images, in-flight calibration images and navigation images from the start of Saturn approach science through the end of mission. These are identified with the following DATA_SET_NAME, DATA_SET_ID and STANDARD_DATA_PRODUCT_ID:

CASSINI ORBITER SATURN ISSNA/ISSWA 2 EDR VERSION 1.0
CO-S-ISSNA/ISSWA-2-EDR-V1.0
ISS_SEDR

3. The calibration files, including calibration data, ground calibration image files, documentation, calibration software and sample calibrated images. These are identified with the following DATA_SET_NAME, DATA_SET_ID and STANDARD_DATA_PRODUCT_ID:

CASSINI ORBITER CALIBRATION ISSNA/ISSWA 2 EDR VERSION 1.0
CO-CAL-ISSNA/ISSWA-2-EDR-V1.0
ISS_CAL

The VOLUME_ID for each archive volume is COISS_xxxx, where the first x = 1 for Jupiter, 2 for Saturn, 3 for cartographic maps, and 0 for calibration, and where the next xxx is the sequential numbering of the volume starting with 001. The VOLUME_NAME of the archive volume depends on the dataset type and contents as follows:

CASSINI ISS EARTH/VENUS/JUPITER EDR SCLK xxxxxxxxxxxx to SCLK yyyyyyyyyy
CASSINI ISS SATURN EDR SCLK xxxxxxxxxxxx to yyyyyyyyyy
CASSINI ISS CALIBRATION FILES

where xxxxxxxxxxxx and yyyyyyyyyy are the 10-digit start and stop SCLK counts of the volume contents.

3.3 ISS VICAR Image Format

Each individual ISS image is contained in a file called an Experimental Data Record (EDR), which is archived in a format called VICAR (Video Image Communication And Retrieval). VICAR is an entire system of software, formats, and procedures for image storage and processing and was developed and is maintained by JPL's Multi-mission Instrument Processing Laboratory (MIPL). A full explanation of VICAR, its standards, software and reference information can be found at the website:

<http://www-mipl.jpl.nasa.gov/vicar/>

Information on tools for visualizing VICAR images can be found there as well. For example, the PDS-provided NASAview tool can be downloaded from the PDS site (<http://pds.jpl.nasa.gov>) and used to view the raw images. Code for reading VICAR images are also available for the Interactive Data Language (IDL); see for example the links at <http://atmos.nmsu.edu/Jupiter/jupiter.html>.

The file name of the image data file consists of a string that incorporates information about the instrument name, spacecraft clock and version number. The instrument name is required because the spacecraft clock is not unique for simultaneous exposures. The version number is required because the same image may be built multiple times due to multiple downlinks, or multiple Telemetry Data System (TDS) queries, etc. The following naming conventions are followed:

Image file = <camera><SCLK time>_<version>.IMG

Where:

camera = 1-character instrument identifier (N=NAC, W=WAC)

SCLK time = 10-digit value of spacecraft clock at time of shutter close

version = version number of the file

So, for example, an image file named W1832898283_4.IMG would indicate a the fourth version of a Wide Angle Camera image taken at SCLK time1832898283. (The corresponding detached label file

follows the same naming convention as above except with ".LBL" as the filename extension. Example: W1832898283_4.LBL)

Each image data file contains several fixed-length data records. These are: the ASCII VICAR Label (also known as the “attached label” or simply "image header"), the Binary Label Header (or "Binary Telemetry Header"), and the Image Line Records, which are comprised of the Binary Line Prefix plus the actual pixel data. All of these are briefly described in the paragraphs below. For more complete information about the format and content of the image data products, see the Cassini ISS EDR Software Interface Specification document (edrsis.txt, edrsis.pdf) found in the document/ directory of these volumes. These image files are reconstructed from the best available telemetry data and line-filled where necessary to produce the most complete image records possible.

The ASCII VICAR label is included to facilitate processing and allow easy validation and traceability of the image products from version to version throughout the downlink process. This image header consists of a set of ASCII "KEYWORD=VALUE" pairs describing the important characteristics of the image. It is designed to be human-readable so that label items can be easily extracted in order to guide automated processing procedures, or to annotate products derived from the image, such as plots.

The VICAR Label contains System items which contain structural and format information about the data file, as well as several types of Property items which describe the various instrument settings and other observation and image characteristics. The VICAR Property label types used by Cassini ISS are: COMMAND, COMPRESSION, IDENTIFICATION, IMAGE, INSTRUMENT and TELEMETRY, and these are identified in the label by a preceding PROPERTY keyword (e.g. PROPERTY='INSTRUMENT'). Individual property keywords, their sources and valid values are given in the EDR SIS.

The Binary Label Header (also known as the Binary Telemetry Header) contains 60 bytes worth of machine-readable information about the image as a whole, padded with zeros to the end of the image

record length. Many of these items are in the VICAR Label as well, but non-VICAR sites may ignore the VICAR Label and use the Binary Telemetry Header to construct their own human-readable label. Items in this header are copied directly from the Extended ISS Science header returned in telemetry.

There is one Line Record for each image line, comprised of a 24-byte Binary Line Prefix followed by the 256, 512, or 1024 pixels' worth of 8- or 16-bit pixel data for that line. The Prefix contains information about the image line derived from telemetry: Line Number, the Extended and Overclocked pixel values for that line (see below), and values indicating the starting and ending locations of the valid pixel data. Because the pixels in a given line may come from up to four data packets, a line may consist of segments of good data accompanied by segments of zero-filled missing data. The Prefix specifies the location of the beginning and end of up to two of these “good” line segments.

Note that for Lossy compressed images, the data are not associated with lines, so there is no way to associate a given record with a line number. In this case, the Binary Line Prefix contains information extracted from the information received in the last compression block.

The detector system includes an un-illuminated region eight samples wide - the “extended pixel” region - extending into the negative sample direction in the serial register. These pixels get read out first. Moreover, once an entire row of 1024 pixels is read into the serial register and out to the signal chain, the readout continues for eight more clock cycles, or “overclocked pixels,” to provide a measure of the offset bias, the DN value that corresponds to zero signal level. The extended pixel region and the overclocked pixels in principle provide two independent measures of offset bias and a sample of the horizontal banding pattern that may be used to remove the pattern in images lacking dark sky. (A discussion of the horizontal banding problem can be found in Section 4, and in Porco et. al., 2004.)

The Binary Line Prefix does not contain individual entries for each overclocked or extended

pixel in a line, but rather sums of several values, the specifics of which depends on the compression mode and flight software version. The overclocked pixels are actually recorded as two separate values - the First Overclocked Pixel Sum represents the sum of the first two overclocked pixels (or just one if not in 1x1 summation mode) which are clocked out *before* the line, and the Second Overclocked Pixel Sum represents the sum of the last six overclocked pixels in 1x1 mode, last three in 2x2 mode, or last one in 4x4 mode, clocked out *after* the line. Similarly, the Extended Pixel Sum is a sum of eight extended pixels in 1x1 mode, four in 2x2 mode, and two in 4x4 mode. For lossy-compressed images, sums are returned only for the lines in the last compression block, with zeros filling in otherwise.

The image label also contains two single-value keywords representing averages of the overclocked and extended pixel arrays – these are the BIAS_STRIP_MEAN and DARK_STRIP_MEAN, respectively. For these averages, the first and last lines in the binary line prefix are ignored.

In the NAC, the extended region of the readout register, and the first 13 columns into the serial register are corrupted by a grounding problem with the epoxy that bonds the pure silicon layer to the substrate. This causes spurious swings in the voltage during the clocking out of the CCD data into the signal chain. Consequently, these columns of CCD data are unreliable, and the NAC's extended pixel region cannot be used to monitor the camera's bias or noise state.

The overclocked and extended pixels, and the BIAS/DARK_STRIP_MEAN values that are derived from them, have changed since launch due to flight software upgrades. Table A.9 in the appendix provides a complete list of these changes.

3.4 Reading ISS Images

There are a number of software packages capable of reading and displaying ISS VICAR files, including public domain software like NASAView and ISIS (Integrated System for Imagers and Spectrometers), both of which are available for download at the PDS software download site:

http://pds.nasa.gov/tools/release/software_download.cfm

Additional options may be found at the Planetary Rings Node ISS software page:

<http://pds-rings.seti.org/cassini/iss/software.html>

For the examples in this manual we will assume the use of IDL (Interactive Data Language) as this is the language used for the ISS Calibration pipeline CISSCAL. To simply read a VICAR file into an IDL array for display, one can use the `read_vicar.pro` routine found here:

<http://atmos.nmsu.edu/Jupiter/jupiter.html>

or any number of alternative “read vicar” routines easily accessible via a google search. CISSCAL itself also contains basic image display options.

3.5 PDS Label

Every ISS image archived with PDS has an associated detached PDS label file, which is identically-named, but with an “.LBL” file extension in place of “.IMG”. These files, described completely in the EDR SIS (`edrsis.txt`, `edrsis.pdf`; located in the `document/` directory of the ISS archive volumes) contain identical keyword information to that which is contained in the attached VICAR Labels. In addition, they contain pointers to and information about the data objects within the image

file – specifically, the IMAGE_HEADER, TELEMETRY_TABLE, LINE_PREFIX_TABLE, and IMAGE. These objects are denoted by a statement of the form:

`^object = location`

in which the caret character ('^', also called a pointer in this context) indicates that the object starts at the given location, i.e. the name of the file containing the object, along with the starting record or byte number. For example:

`^IMAGE = ("N1294562651_1.IMG",3)`

indicates that the IMAGE object begins at record 3 of the file N1294562651_1.IMG, in the same directory as the detached label file.

The detached PDS labels should generally be kept in the same directory alongside their corresponding .IMG files. While the detached labels are not used by CISSCAL for calibration (as all the required information is available in the attached label/image header), they are required by some other software packages, including both ISIS and NASAView.

3.6 ISS Data Quality and Completeness

The quality and completeness of the image data are determined in several phases. Firstly, within MIPL, images are constructed from the raw data stream using automated MIPS-provided VICAR software. Image construction is performed by replaying telemetry from the DSN in multiple passes to obtain the best possible image products. This process may require reconciling between overlapping data streams from multiple DSN stations, with the best available telemetry from each station being used to construct the image data.

Next, verification software is used to generate product and quality reports that detail missing

images and incomplete data from the downlink, and the reason for the discrepancy. These reports are included on the DATA volumes, in the /document/report/ subdirectory. (Note: no product and quality reports were generated for images prior to SCLK 1431917000.) The contents and formatting of these reports are detailed in the dataset.cat file.

Finally, the ISS team routinely compares the images returned against the predicted images from observation planning and uplink process. This is done as part of the ISS team's normal data usage and science analysis. In addition, CICLOPS-generated scripts are run by team members to ensure all images posted by MIPL to the server are indeed received by CICLOPS and are maintained in the ISS archive database.

Keyword values are subject to inaccuracies; usage is cautioned. Keywords that come directly from the image label are included verbatim and are as reliable as the sources of those keywords (i.e. MIPS telemetry processing software using spacecraft and camera commanding file inputs). More crucially, the accuracy of the geometry keywords in the index.tab is dependent on the accuracy of the auto-navigation software as well as the accuracy of the various SPICE kernels used to calculate the keywords.

The quality and completeness of the archive volumes generation process are also determined by the accuracy of the archive generation software written and used by the CICLOPS team. This software identifies which images are to be archived on a given DVD volume and then constructs the archive directory structure with the appropriate image and label files and static documentation. The dynamic information files are updated as needed and reviewed for accuracy. Finally, another CICLOPS-generated script is then run on the volume to check for obvious mistakes or omissions, before it is burned to a DVD and delivered, physically and electronically, to PDS.

In order to ensure PDS-compliant products, the archive volumes are validated by a collaborative effort between the ISS/CICLOPS team, the Imaging and Central Nodes of the PDS, and non-Cassini

imaging scientists. Validation is considered to have two aspects: 1) scientific usability and 2) technical compliance to PDS standards. PDS standard compliance checks are performed by PDS using their own validation software, as well as by CICLOPS-developed verification tools run prior to delivery to the PDS Imaging Node.

Scientific usability is assessed through the ISS science team's normal and routine use of the ISS datasets in their science analysis. Additionally, imaging scientists not associated with the Cassini project participate in the archive volume peer review process where they verify the content of the dataset, the completeness of the documentation, and the scientific validity (i.e., the integrity and usability) of the datasets.

Descriptions of the most egregious data anomalies are provided below. More information regarding anomalous data, erroneous keywords, and other errata items concerning the ISS archive volumes can be found in the `errata.txt` and the “Confidence Level Note” in the `dataset.cat`.

Truncated Images

There are several causes of image truncation in the ISS. Packets may be lost in transmission, either due to malfunction, inclement weather on the ground or some other interference. Whole and partial images may be lost due to data policing, when the allocated data volume for the observing segment is exceeded. Another type of truncation has to do with the compressibility of data in losslessly-compressed images, and results in the loss of data from the ends of alternating image lines. For more information about this issue, see the relevant discussion in Section 2.3, and the example in Figure 4.

The last type of image truncation is caused by an insufficient image readout window, which may be intentional, or the unintentional result of a poor compression estimate during the observation planning process. This type of truncation results in the loss of an entire section of the image that was

unable to be read out before the end of the allotted window.

Take for example an uncompressed, un-summed, 12-bit NAC image. The camera generates 2277 packets, and at a data telemetry rate of 24 packets per second, takes about 95 seconds to read out. (See the readout index and telemetry tables in Appendix A.) If the observation designer had chosen a readout index from 4 through 7, corresponding to 50 seconds, then the readout would only be half-finished when the readout window completes, and the resulting image would be partial. To avoid truncation in this case, a readout index between 0 and 3 is required. It is easy to see how the readout time for uncompressed images can significantly limit the number of exposures that can be taken in a given time period. In some cases, the observation designer may want to image more quickly than what is possible with full image readout, and accept some amount of image truncation.

It gets more complicated when images are read out in a compressed mode, since the amount of data to be transmitted from camera to spacecraft depends on how well it compresses, and that isn't known explicitly until the image is acquired. Take for example a 1x1, 12-bit, lossless image with an expected compression ratio of 5:1. One would expect 461 packets and a readout time of 19.5 seconds, so ISSPT chooses readout index 8 (28 seconds for a telemetry rate of 24 packets/sec). Now, say the image had higher entropy than expected and only compresses at a ratio of 3:1. The actual number of packets is 764 requiring at least 32 seconds of readout time, but the readout will stop after 28 seconds and the last 1/8 of the image will be lost.

Lossy Compression Camera Bug Anomaly

An anomaly in the NAC and WAC camera software (FLIGHT_SOFTWARE_VERSION_ID = 1.3) was discovered in April of 2004. This machine error is caused by the retrieval of extended and overclocked pixels in images taken in LOSSY compression mode. A fix was executed in September of 2004 to correct the problem. A significant number of images were lost due to this bug between the

SCLK times 1462417483 and 1481784349. These missing images are noted in the quality reports with the Instrument Surprise Anomaly (ISA) number listed in the 'REASON' column. Cassini ISA reports Z83951, Z83931 and Z84199 were filed to document the problem, although these will be accessible only to operations personnel for the duration of the mission.

Double Bit Error Anomaly

A bad DRAM submodule on the SSR has been corrupting approximately 1% of ISS data since 2006 DOY 338. These “double bit errors” affect words in the frame headers and the header and data portions of ISS image packets. This may cause corrupt values in the ISS standard and extended header (e.g. invalid temperatures or camera parameters resulting in predict mismatches which cause packets to be thrown out), corrupt line headers which cause a packet to be thrown out, or corrupt data values which can cause either bad pixels in uncompressed images, errors in the compressed stream causing line pairs in lossless images to show errors, or truncated lossy images occurring at the first dropped packet due to mismatch errors. Packets which would otherwise have correct ISS header data (standard, extended and line) but have invalid frame (spacecraft) data cause lost packets, which are thrown out before our processing pipeline sees them and will show up as missing data in uncompressed and lossless images or truncated lossy images. See images N1543791633_3.IMG and W1543791633_3.IMG for examples of images corrupted by the double bit error anomaly.

The latest report detailing periods of suspect data due to double bit errors is pasted at the end of the errata.txt file. The best way to query for these images is to search for instances where `MISSING_PACKET_FLAG = 'YES'`.

A flight software fix is planned for the future which will skip over the bad submodules on the solid state recorder. Data affected by the anomaly may or may not be reconstructed in the future. The

errata.txt file will be updated with more information as it becomes available.

NAC Haze Anomaly of 2001

In May 2001 (Day 150), in NAC images taken of the Pleiades, a diffuse circular halo appeared around the central peak of the image of Maia; WAC images were not similarly degraded. The apparent cause of this anomaly was the resumption of normally scheduled decontamination cycles after a 13-month hiatus. Additional conservative decontamination cycles were performed and the haze disappeared leaving the point response function of the NAC within pre-anomaly limits. Subsequent testing has shown no apparent change in the instrument response before and after the decontamination. For more detailed information on this anomaly, see the document titled “Cassini Camera Contamination Anomaly: Experiences and Lessons Learned” obtainable from the JPL Technical Report Server here: <http://hdl.handle.net/2014/40797>.

3.7 Archive Volume Organization and Format

Archive Volume DVD Format

The PDS ISS archive is organized into volumes according to the number of images that can fit on a standard one-sided 4.7 GB DVD, and the directory structure and file contents of each volume reflects that of an actual physical DVD delivered to PDS. In order to ensure that a wide variety of computer systems can access the data, the DVDs have been formatted according to the ISO-9660 Level 2 Interchange Standard. (For further information, refer to the ISO-9660 Standard Document: RF#ISO 9660-1988, April 15, 1988.) The volumes have both ISO and UDF file systems. With the exception of

the EDR product files (*.IMG, *.LBL), all file names on the volumes should be lower-case. Filename case may not be preserved if the host computer system reads the ISO file system instead of the UDF file system; though most computers should default to read the UDF file system.

A complete description of the archive volume directory structure is contained in the aareadme.txt and archsis.pdf, but the best and easiest way to familiarize oneself with the contents of the archive is through a web browser, starting with any of the volume on the PDS Imaging Node's ISS page:

<http://pds-imaging.jpl.nasa.gov/volumes/iss.html>

With the exception of the aareadme.txt, the errata.txt and the voldesc.cat, all files on the ISS archive volumes are organized into PDS-standard subdirectories below the top-level directory. These are the catalog/, data/, document/, index/ and label/ directories on the DATA volumes and first ten CALIBRATION volumes, and all of these plus the extras/ and calib/ directories (containing the calibration software and support files, respectively) on the eleventh (COISS_0011) CALIBRATION volume. Each subdirectory contains an ASCII text “info” file (e.g. catinfo.txt, labinfo.txt, etc.) describing the contents of that directory. Furthermore, most ASCII text files (.cat, .tab, or .txt extensions) contain attached PDS labels with additional descriptive information.

The organization of the image data on the data volumes is straightforward. Image files and their PDS labels are contained in the data/ directory, organized into subdirectories named according to the first and last IMAGE_NUMBER contained within, i.e.:

<IMAGE_NUMBER_start>_<IMAGE_NUMBER_end>

Each subdirectory contains 128 image files and 128 label files, for a maximum of 256 files per subdirectory (per PDS requirements). NAC and WAC image files are situated side-by-side in the same subdirectories, although, because they are organized by the IMAGE_MID_TIME (as opposed to the time of shutter close which defines the IMAGE_NUMBER) in some cases it is possible that two

images in a BOTSIM will end up in separate subdirectories.

The number of images per directory and per volume varies, as does the amount of data in terms of bytes. This is because, starting in October of 2004, new data volumes have been released on a 3-month schedule, and the amount of data taken in any 3-month time period varies considerably. Typically, for one of these deliveries, there will be two volumes filled almost to capacity, and a third that is only partially filled. Roughly-speaking, approximately 3,000 to 5,000 images can be found on each archive DVD volume.

Calibration Volume Organization

As discussed above, the first ten volumes of the calibration data set contain the pre-flight ground calibration images. Volume eleven contains the collection of calibration data files, calibration software files, sample calibrated images and related documentation.

The ground calibration files were originally distributed on a collection of CD-ROMs by the Instrument Operations team. They were intended to be PDS-compliant when produced. The Imaging Node later converted the CD-ROMs to DVDs for inclusion in this ISS archive collection.

The ISS archive contains an extras/ directory that is included only on the eleventh Calibration volume. This directory contains the source code for the Cassini ISS Calibration (CISSCAL) software. This software, developed by the Cassini Imaging team, allows the user to radiometrically and geometrically process the EDR-level images into higher level calibrated images.

CISSCAL was developed using the Interactive Data Language (IDL). No compiled executables are supplied; IDL Version 5.5 or later is required to compile and run the code. Note that, in the case that your computer system reads the ISO filesystem (instead of the UDF filesystem) of the calibration DVD volume, filenames may display as uppercase instead of the default lowercase. This will break certain filename references in the CISSCAL software. To get around this issue, the entire contents of

the cisscal/ subdirectory have also been provided as a gzipped TAR archive.

The CISSCAL manual, cisscal_manual.tex, is located in the document directory on the eleventh Calibration volume, although this Data User's Guide makes that document largely obsolete.

The calib/ directory on the eleventh Calibration archive volume contains the calibration data files (sometimes called calibration “support” files) used by CISSCAL for processing the raw EDR images into higher-level products. The calibration data files range in format from text files (filter transmission functions, effective wavelengths, etc.) to VICAR image files (bright-dark pixel pair maps, flatfields, etc.), to Tagged Image File Format (TIFF) images and assorted binary-format data files.

The calib/ directory is formatted in such a way as to be compatible with CISSCAL; users who wish to use CISSCAL will want to copy the entire calib/ directory intact to a location on their local filesystem where they have write privileges. As with the CISSCAL subdirectory, the contents of the calib/ subdirectory are also provided as a gzipped TAR archive to avoid filename case problems that may arise when reading the DVD on some computer systems.

Both the calibration software and calibration data files will be updated throughout the mission; this may include newly-generated data files.

The voluminous ISS Ground Calibration Report can be found on all of the calibration volumes in the /document/report/ sub-directory as a hypertext (HTML) file which is easily navigated with a web browser. See Section 4 for more about ISS image calibration.

Document File Formats

ISS archive document files can be found in one or more of the formats designated by the following file suffixes: .txt, .tab, .cat (ASCII text files), .htm (Hypertext Markup Language), .tex (LaTeX/TeX typesetting language), and/or .pdf (Portable Document Format).

According to PDS requirements, a human-readable ASCII text version (with “.txt” filename

extension) must be included for each document. These text files have line lengths restricted to 78 characters or fewer in order to accommodate printing and display on standard devices. Each line is terminated by the two-character carriage-return/linefeed sequence, <CR><LF> (ASCII decimal character codes 13 and 10, respectively), for a maximum total line length of 80 characters.

The .txt document files may be accompanied by corresponding .pdf document files in cases where documents contain formatting and figures could not easily be rendered as ASCII text. Portable Document Format (PDF) is a proprietary format of Adobe Systems Incorporated that is frequently used for distributing documents.

LaTeX is a high-quality typesetting system, with features designed for the production of technical and scientific documentation. LaTeX files have relatively little markup embedded in the text and are generally considered human-readable and may, therefore, be used to satisfy the ASCII text version requirement. One exception may be tables within the document, which will not appear properly typeset unless the LaTeX file is first compiled and converted to a different format. LaTeX is free, and is currently developed and maintained by LaTeX3 Project. Information about the system and various conversion software can be found at their current website: <http://www.latex-project.org>.

Files written in hypertext markup language (HTML, .htm suffix) contain ASCII text plus commands that enable viewing of the document in a web browser. The hypertext file may be accompanied by ancillary files such as images and style sheets that are incorporated into the document by the browser.

Tabular-format files (.tab suffix) can be found in the index directory and in several of the calib/ subdirectories. Tabular files are ASCII text files formatted for direct reading into many database management systems on various computers. All tabular files are described by either detached or attached PDS labels. The index.tab file, containing the complete set of PDS image keywords, is described in more detail in the following subsection.

Catalog files (.cat suffix) exist in the catalog directory. They are text files formatted in an object-oriented structure consisting of sets of “KEYWORD=VALUE” declarations. Each line is restricted to 72 characters or fewer, and is terminated by the two-character carriage-return/linefeed sequence <CR><LF> to accommodate PDS data ingestion requirements set forth by their internal catalogs and databases.

3.8 The Index.tab File

The image index table file, index.tab, contains the most complete set of keyword information for each image on the volume. Some of this information comes directly from the image label produced by IO - all of the camera and instrument settings, for example; some comes from the archive generation process, like the DATA_SET_NAME and INSTRUMENT_HOST_ID; and the remaining keywords come from the Autonav software (discussed in Section 3.1 above) which calculates the many geometric quantities and target information keywords listed in Table 5.

Each line in the index.tab is a record containing all the keywords for a particular image on the volume. Fields in a record are delimited by commas, and non-numeric (text string) fields are enclosed in double quotation marks ("), left-justified, and padded with spaces to keep the column widths even. Numeric fields are right-justified. Multi-valued fields are enclosed in brackets and each item in the field is separated by a comma. The records are of fixed length, and the last two bytes of each record contain the ASCII carriage-return/line feed character sequence, <CR><LF>. This allows the index.tab to be treated as a fixed-length record file on computers that support this file type and as a text file with embedded line delimiters on those that don't.

Perhaps the easiest way to view the index.tab file and search its contents is to simply treat it as a CSV (“comma-separated values”) format text file and read it into a spreadsheet program such as

Microsoft Excel or OpenOffice.org Spreadsheet. For the latter, this requires first renaming the file to give it a .csv file extension before opening it. When the “Text Import” dialog box opens up, set the field delimiter set to comma and the text delimiter set to double quotes (“”). Once the index file has been read in to the spreadsheet, its contents can be searched and sorted according to any desired combination of keywords.

The associated detached label file, index.lbl, provides a complete file specification, including the name, data type, start byte, number of bytes, and format of every PDS keyword in the index.tab. Additionally, the full list of keywords with descriptions, valid values and other useful information is provided in the archsis.pdf.

3.9 Cartographic Maps

In addition to the raw EDR image data, calibration files, and supporting metadata, the ISS team also archives higher-level data products with the PDS in the form of cartographic maps of Saturn's icy satellites. These are stored on their own archive volumes (VOLUME_ID = “COISS_3XXX”). As of this writing, seven of the medium-sized icy satellites are represented; ordered by their volume number (XXX, above) they are: Phoebe (001), Enceladus (002), Dione (003), Tethys (004), Iapetus (005), Mimas (006) and Rhea (007). These volumes can be accessed directly from the PDS Imaging Node here:

<http://pds-imaging.jpl.nasa.gov/volumes/carto.html>

by clicking on “ISS Cartographic Maps.”

From the cartographic map volume aareadme.txt: “One of the objectives of the ISS team is to obtain global coverage for all medium-sized icy satellites with a resolution better than 1 km/pixel and

obtain high-resolution images. This goal is being achieved with image sequences obtained during close flybys supplemented by images from greater distances to complete the coverage. Close flybys of all medium-sized satellites except Mimas are planned during the nominal mission of the Cassini spacecraft[...] The cartographic maps were created from images taken by Cassini with a very small number of images from Voyager-1 and Voyager-2 used to fill in any gaps. Only images taken with CL1, CL2, and GRN filters were used[...] Imaging of the medium-sized icy satellites is ongoing and will continue until the end of the Cassini mission, making it possible to improve the image mosaics during the tour.” For more information about ISS cartographic maps, consult the aareadme.txt and other supporting documentation on the cartographic map archive volumes.

3.10 Accessing ISS Data Online

Most users will want to access ISS data electronically over the internet. There are a few different options for doing this. Direct access to the PDS ISS archive volumes is available from the PDS imaging node here:

<http://pds-imaging.jpl.nasa.gov/volumes/iss.html>

(You can also get there from the Imaging Node homepage, by clicking on “Data Volumes Index” and “ISS” under the “By Mission” heading.) PDS provides links to gzipped tar archive files containing complete volumes, as well as links to servers containing the entire directory structure of the archive volumes, with the VOLUME_ID as the name of the root directory; for example, the following URL:

http://pds-imaging.jpl.nasa.gov/data/cassini/cassini_orbiter/coiss_2001/

directs the user to the root directory of the COISS_2001 volume, which is to say the first volume of the Saturn EDR data set.

Accessing ISS data in this way is most useful for those who already know which images they are looking for, or are limiting their search to a particular time range represented by a particular data volume or range of data volumes. If one would like to search the entire data set for images before downloading them, PDS provides the “Planetary Image Atlas” search tool:

<http://pds-imaging.jpl.nasa.gov/search/search.html>

This tool gives users the ability to query, either singly or in combination, most of the ISS PDS keywords contained in the index.tab file, including the target geometry keywords generated by Autonav.

To begin a query for ISS data, first set the Mission to “CASSINI” and the Instrument to “ISS” using the menus in the left side panel. This will ensure that only results for the ISS are returned, rather than for all imaging data archived with PDS.

The search interface is divided up between five tabs, with keywords grouped roughly according to type. The major exception to this organization scheme is that many of the most useful or commonly-queried keywords – including FILTER_NAME, TARGET_NAME, INSTRUMENT_ID, and INCIDENCE/EMISSION/PHASE_ANGLE - are found under the “QuickSearch” tab. All other keywords are split up between the “Product,” “Geometry,” “Instrument,” and “Time” tabs. The interface is designed such that keywords with freely-varying numerical values (i.e. not text or indexed values) are queried by filling in “minimum” and “maximum” fields, both of which must be specified for the search to work. Keywords with arbitrary text values are queried via text fields, and the queried text need not be the full keyword value, but rather can be any sub-string. Keywords with indexed values are queried either by toggle button (e.g. INSTRUMENT_ID is either “ISSNA” or “ISSWA,” ANTIBLOOMING_STATE_FLAG is either “ON” or “OFF”) or via scroll-down menus containing the valid keyword values. The scroll-down menus also allow the user to select more than one search value from the list, by holding down the shift key while clicking. (To de-select the last selected item in a

scroll-down menu, one may need to hold down the control or option key while clicking, depending on the browser and operating system.)

Particular care should be taken with the FILTER_NAME keyword value, as it is not explicitly clear which filter combinations belong to which camera. For instance, a user who searches for INSTRUMENT_ID="ISSNA" and FILTER_NAME="CL1,BL1" may be surprised to find zero returned results, when in fact, because of the different layouts of the NAC and WAC filter wheels, the correct name for this filter combination in the NAC is "BL1,CL2."

Note also that the search interface returns results for both EDR (raw) image data, and higher-level RDR products; specifically, in this case, cartographic maps of some icy satellites. To limit the search to EDR images only, click the "EDR" checkbox next to PRODUCT_TYPE in the QuickSearch tab.

Once query selections have been made, click "Get Count" on the left side-panel to find the number of images returned by the search, and "Get Results" to view the results page. Due to its support of multi-mission, multi-instrument functionality, the search results page has a few idiosyncrasies that should be noted. Chief among these is the way the keyword columns are displayed – queried keywords are not included in the results columns by default, but rather must be explicitly selected after the fact using the "SELECT PARAMETERS..." scroll-down menu on the right side of the screen. Furthermore, these selections are lost each time the search parameters are changed, although the search parameters themselves, luckily, are not. The PRODUCT_ID, MISSION, and INSTRUMENT columns are always displayed for all searches.

Another quirk of the Planetary Image Atlas results page is the way it deals with searches that return large numbers of images. By default, only the first 500 matches to a search are returned – the text centered above the results columns will say, for example, "500 of 1632 Products." To ensure that all images are included in the search results, click on the "Get More" button next to the page numbers

at the upper left to add the next 500 images to the results. (A search returning 2000 images would require clicking on “Get More” three times in order to display the entire list.) To control the number of search results displayed per page, use the “Results Per Page” drop-down menu at upper right. Note that this is a separate value than the total number of results returned, which can only be changed by clicking “Get More.”

From the results page, there are several possible actions one may take. For each search result, three icons are provided in the results table which link respectively to the image file, to the PDS image label, and to the root directory of the archive volume on which it the image can be found. Batch download options are also available. Use the “Select” column to select multiple results from the results table, and then under “Download Products” in the right sidebar, select either the “WGET” or “ZIP” download option. The “ZIP” option simply returns a compressed .zip file. This requires first creating a .zip archive before the download can begin, which can be quite time-consuming when downloading large numbers of images. For many users, wget will be the better option. Selecting “WGET” returns a script file – a simple text file, which by default is named “atlas_wget_script” – that can be run as an executable from the Linux/UNIX command line, and will automatically download the selected image files to the directory from which it is run. To use the wget script, first save the file to the desired directory, open a terminal window and cd to that directory. Then at the command line, type:

```
chmod u+x atlas_wget_script
./atlas_wget_script
```

Assuming you have the wget executable in your PATH system variable (standard on most installations), your download should start immediately. Options are available for running wget on Windows as well.

Users may also find it useful to download comma- or tab-delimited text files containing the search results table (and all selected keyword columns). These are obtained by selecting the “CSV” or “TAB” options respectively, under “Download Report” on the right sidebar.

Further suggestions and examples for using the PDS Planetary Image Atlas search tool can be found in Section 5 of this document.

4. ISS Image Calibration

4.1 Introduction to ISS Calibration

The goal of this chapter is to describe, or summarize where necessary, all of the information required for calibrating Cassini ISS images for scientific purposes. We begin by introducing the theoretical basis for radiometry that underpins the CISSCAL software pipeline, and the basic order of operations for the calibration process. Next we provide a comprehensive summary of the calibration data that characterizes the detailed behavior of the instrument. Some of these data are derived from in-flight observations, but the majority is the result of pre-launch calibration testing conducted at JPL's Environmental Test Laboratory. These tests were performed at the individual component level as well as the subsystem level under thermal vacuum conditions, and many of the camera parameters derived from the subsequent image analysis are stored in data files that are provided with the ISS archive and used by CISSCAL for calibration. This discussion will also cover the various in-flight calibration observations that are scheduled throughout the mission, the results of which are periodically incorporated into updates to the calibration data files. Finally, we describe the CISSCAL calibration software package along with strategies for its use.

The subsections that follow will summarize and attempt to clarify much of the calibration documentation already included with the ISS Calibration volumes – namely, the `theoretical_basis.pdf`, the ground calibration report, the `in_flight.pdf`, and the `cisscal3_6_manual.pdf` included with the CISSCAL distribution (see Table 1 in the Introduction). This document intends to replace and supercede the CISSCAL user's manual (`cisscal3_6_manual.pdf`) currently included with the CISSCAL distribution, and will borrow text from that document where appropriate.

Some knowledge of astronomical charge-coupled device (CCD) detectors is assumed on the

part of the reader, including basic vocabulary and familiarity with fundamental concepts like bias level, dark current, and flatfield. The CISSCAL software pipeline is written in the Interactive Data Language (IDL) and that package is required in order to run the specific algorithms discussed here. Calibration of the ISS is a work in progress, and the ultimate goal of this guide is to provide a savvy user with enough information about the instrument, its capabilities and its limitations that she or he will be able to pick up where the CISSCAL developers left off, and add to or improve on the existing algorithms where called for in the course of scientific research.

4.2 Theoretical Basis of Radiometric Calibration

In order to properly understand the steps involved in calibrating the ISS, it is important to first establish a theoretical and mathematical basis from which to proceed. Here we will present a slightly condensed version of the information available in the `theoretical_basis.pdf` document in the `document/` directory of the Calibration volumes, and concentrate on the essential aspects that relate to the CISSCAL implementation and the various camera parameter dependencies that it must take into account.

The goal of calibration is to relate the data number (DN) values recorded at each pixel location in an image to actual physical units of incident intensity, thus making the images suitable for scientific analysis. The essential task is encapsulated in the following equation relating DN to intensity I as a function of image sample number i and line number j :

$$DN(i, j) = \frac{A\Omega t(i)}{g} \int I(i, j, \lambda) T_o(i, j, \lambda) T_1(i, j, \lambda) T_2(i, j, \lambda) Q(i, j, \lambda) d\lambda + DN_o(j) + DN_D(i, j) \quad (1)$$

Here, A is the collecting area of the optical primary, Ω is the solid angle subtended by a pixel, t is the

exposure time (sample-dependent due to the shutter motion), g is the gain in electrons/DN, I is the incident intensity, DN_0 is the bias level, DN_D is the dark current (about which, see below), Q is the detector quantum efficiency in electrons/photon, and T_o , T_1 and T_2 are the wavelength-dependent fractional optics, filter 1, and filter 2 transmissions, respectively. The optics area and solid angle are known, and the properties of the gain, filter and optics transmission, the sample-dependency of the exposure time due to the shutter mechanism, and the detector quantum efficiency can all be measured separately in the laboratory. The bias level DN_0 varies with line number due to low-frequency variation in the bias voltage during readout (aka “2 Hz noise”) and must be dealt with on an image-by-image basis by isolating the 2 Hz signal either from the overclocked pixels or from the image itself.

In the ISS, the quantity DN_D is dominated almost entirely, not by dark current in the traditional sense – the detector operating temperature of approximately -90° C is enough to ensure very low dark current buildup except at the very highest exposure times – but rather by an effect known as “residual bulk image,” or RBI. Residual bulk image is the result of charge from a pixel in the previous image becoming trapped in the silicon layer of the CCD chip and leaking back into the potential well of the subsequent image. To help control the effects of RBI, ISS exposures are by default immediately preceded by a “light flood and erase” step, whereby the detector is flooded with light from infrared LEDs to many times the pixel saturation level, and then read out. This procedure fills all of the electron traps in the silicon, ensuring a repeatable starting condition for each image and eliminating any RBI spatial dependency from exposure to exposure. It is primarily the RBI from this light flood event that constitutes the quantity DN_D that must be removed in calibration. Furthermore, because the dark current/RBI is dependent on the readout behavior of the CCD, and because the readout in turn depends in complex fashion on the exact combination of camera settings, the conventional measure of DN_D - averaging a set of dark images obtained with the same camera settings, preferably acquired at the same time as the observation - is not feasible. Instead, a modeling approach is used. This will be described

further in the discussion of the dark removal implementation in Section 4.4.

The quantities inside the integral in Equation 1 all vary with wavelength as well as spatially across the CCD detector. Rather than having to treat the spatial variation in the optics, filters and detector quantum efficiency separately, we can define $FF_{f1,f2}(i,j)$ to represent the unity-normalized spatial dependency of the entire system, thus allowing us to pull it out of the wavelength integral. This quantity, called the flatfield, is dependent on filter combination (f1, f2), and is obtainable through pre-flight or in-flight observations of flat, featureless targets.

Finally we introduce the quantity $C_{f1,f2}$, a filter-dependent absolute flux correction factor that will account for any discrepancies in ground-based measurements of the camera components. The equation above can then be rewritten as:

$$DN(i, j) = \frac{A\Omega t(i)FF_{f1,f2}(i, j)C_{f1,f2}}{g} \int I(\lambda)T_o(\lambda)T_1(\lambda)T_2(\lambda)Q(\lambda)d\lambda + DN_0(j) + DN_D(i, j) \quad (2)$$

The calibration process must also take into account the fact that the gain value g deviates somewhat from perfect linearity over the dynamic range of the instrument. This deviation is significant for the lowest gain state (gain index 0) but is less than 1% over most of the DN range for the other three gain states.

Another adjustment to the DN values is necessitated by a phenomenon known as uneven bit weighting, whereby certain DN levels are underpopulated, and others overpopulated, relative to what would be expected from a simple linear correspondence with increasing flux. This effect is caused by errors in the analog-to-digital conversion process, and can be thought of as small discrepancies in the effective widths of the DN “bins” from 0 to 4095. It is dependent on gain as well as detector temperature. Bit weight correction is performed by mapping the image values to slightly offset values derived from their effective bin widths in order to compensate for the uneven distribution.

ISS images suffer from various sources of noise which must be removed during calibration. One

of these is the so-called “2 Hz noise,” mentioned above, which is caused by an unknown source on the spacecraft, and produces a line-dependent banding pattern in un-summed images. This noise pattern actually appears to be a superposition of two or more frequencies that are out of phase with one other. It can be modeled from the overclocked pixel array, or - in shutter-inhibited images, or those containing a significant amount of dark sky across the entire vertical extent of the field of view – from the background image pixels themselves.

Another source of noise is particular to un-summed images with anti-blooming mode set to ON. In this mode, which is designed to reduce charge overflow into adjacent pixels that can result from full-well saturation, electrons may be removed from certain pixels and trapped in the next adjacent pixel in the line direction. This results in bright/dark pixel pairs. These pixel pairs are easily identified on an image-by-image basis, and can be replaced by the average of their neighbors in the sample direction.

Because the camera optics effectively convolves the incident intensity I with the system transmission, we can only recover intensity integrated over the filter bandpass, which is to say, averaged over wavelength with a weighting function of $T_0(\lambda)T_1(\lambda)T_2(\lambda)Q(\lambda)$. To get to that point from the initial $DN(i,j)$ values, the following steps are required:

- 1) If the image has been LUT-encoded to 8 bits, use the reverse operation to convert it back to 12 bits;
- 2) Correct for uneven bit weighting (12-bit images only; LUT and lossy images cannot be corrected due to information lost in the encoding and compression processes respectively);
- 3) Subtract bias using overclocked pixels, or `BIAS_STRIP_MEAN` value (which is an average of the overclocked pixels);
- 4) Subtract the 2 Hz noise using either the overclocked pixels or, if possible, the image background values themselves;
- 5) Subtract dark current and RBI using a simulated dark frame taking into account all relevant

camera parameters;

- 6) Correct for bright/dark pixel pairs created when the anti-blooming mode is turned on;
- 7) Apply linearity correction for the appropriate gain state;
- 8) Divide image by unity-normalized flatfield for the appropriate filter combination;
- 9) At this point, the image is still in DN units; multiply by gain g to convert to electrons, and then divide by the optical area A , pixel solid angle Ω , and exposure time $t(i)$ which incorporates any constant offset from the commanded time, as well as any exposure variability across the frame;
- 10) Convert to intensity units ($\text{photons cm}^{-2} \text{s}^{-1} \text{nm}^{-1} \text{ster}^{-1}$) by dividing by the integral over wavelength of $T_o(\lambda)T_1(\lambda)T_2(\lambda)Q(\lambda)$.
- 11) Divide by the filter-appropriate radiometric correction factor $C_{f1,f2}$.

This completes the radiometric calibration process. Optionally, we can add:

- 12) Correct for geometric distortion using distortion coefficients and focal lengths derived in flight for each camera/filter; this mostly affects the corners of WAC images.

In many cases, it is desirable to express intensities in terms of I/F , that is, normalized to the incident solar flux F . To calculate I/F , divide the image values resulting from the steps above by the quantity.

$$F = \frac{\int F_1(\lambda)T_o(\lambda)T_1(\lambda)T_2(\lambda)Q(\lambda)d\lambda}{\pi D^2 \int T_o(\lambda)T_1(\lambda)T_2(\lambda)Q(\lambda)d\lambda} \quad (3)$$

where D is the distance from the sun to the target body in AU, and F_1 is the solar flux at 1 AU.

The numbered steps outlined above form the basis for the ISS calibration pipeline. For more details about the implementation in CISSCAL, see Section 4.4 below. For more on the derivation of the various quantities, images, and data tables needed for the radiometric calculation, continue on to the following section.

4.3 ISS Calibration Data

This section summarizes the findings of pre-flight and in-flight calibration analyses performed for the ISS. Data files containing the results of many of these analyses are included with the Calibration archive, with any data that is specifically necessary for image calibration being incorporated into the *calib/* subdirectories used by CISSCAL.

The pre-flight testing documented in the ISS “Ground” Calibration Report provides most of the information we have about the detailed behavior of the ISS cameras. Here we will introduce the reader to the results and conclusions of the pre-flight calibration testing as they relate to image calibration, without delving too deeply into the minutia of the process itself. Those who require more detail should consult the source document, which is included on the Calibration volumes in the *document/report/* subdirectory, and is viewable with a web browser via an HTML interface that links to PDF documents containing the individual report chapters. The appendices of the Calibration Report contain much of the information upon which the *calib/* directory files are based.

The ground calibration tests were conducted at JPL's Environmental Test Laboratory prior to launch between January and August of 1996. Individual component testing was conducted prior to subsystem-level tests under thermal vacuum, and many of the analyses performed and presented separately in the Calibration Report – such as the chapters on RBI (component) and dark current (subsystem) - will here be combined for succinctness and clarity. The NAC and WAC were tested separately which ruled out tests of the simultaneous (BOTSIM) imaging mode, with the exception of some limited BOTSIM testing on the assembled spacecraft while in Solar Thermal Vacuum (STV).

Ground-based measurements are complemented by calibration observations conducted in flight. Parameters derived from in-flight calibration include the correction factors that establish the absolute

flux calibration, and the dark current, both of which are monitored over the course of the mission and periodically updated in the calibration data files. Other parameters, such as the shutter offset and gain ratios, are not generally expected to change over time, but seem to have changed from their pre-launch values. In-flight calibration strategies are given in the `in_flight.pdf` document included in the `document/` directory of the Calibration volumes, and the results of in-flight calibrations conducted up through the XM are presented in West et. al. (2010). (This paper will be included on future updates to the COISS_0011 Calibration volume.)

This discussion will be organized by subject matter according to the major components of the instrument: the optics, filters, shutter, and CCD detector. Instrument-level calibrations will also be described, including point-spread function, absolute calibration and polarimetric calibration. Specific references to any archived data files will be included in the relevant subsection.

Optics

The ISS optics are described in Porco et. al. (2004) as well as Section 4.1.1 of the Ground Calibration Report. The NAC optics subassembly is a Ritchey-Cretien configuration comprised of a primary and secondary mirror and two field-flattening lenses. The WAC optics is the Voyager ISS Wide Angle camera flight spare, consisting of a five-element objective lens. The optical train for both cameras also includes the filter wheel/shutter subassembly; the sensor head subassembly, which contains additional field-flattening elements; and finally the CCD window and detector. The effective optical collecting area is 29.32 cm^2 for the WAC and 264.84 cm^2 for the NAC.

The ISS cameras have been designed to maintain their focus without the aid of moving parts, through careful temperature control of the instrument. Due to its long focal length, which makes it particularly susceptible to temperature effects, the NAC is thermally isolated from the rest of the spacecraft, and contains active heating elements to keep the temperature stable to within tight

tolerances all along the optics barrel and between individual components. The WAC has less stringent image quality requirements, and so its temperature is maintained by the remote sensing pallet (RSP) to which it is attached.

Throughput of the ISS camera optics was measured using collimated light sources. Quartz spacers were placed in the filter wheel positions to recreate the CL1/CL2 filter combination, and the effect of these on the throughput was later removed using the measured CL1/CL2 filter transmissions. The resulting measurements have a wavelength resolution of 2.5 nm, covering the entire range of spectral sensitivity (200-1100 nm in the NAC and 390-1100 nm in the WAC) and are recorded in Appendix F of the Calibration Report as well as the na_optics.tab and wa_optics.tab files in the calib/efficiency/ subdirectory of the Calibration archive volume.

Section 4.1.3 of the Calibration Report describes measurements of the camera focal lengths across their fields of view. The NAC focal length was measured to be 2000 ± 4 mm (3σ uncertainty) with no detectable field distortion. The WAC focal length was measured to be 200.22 mm at the paraxial point, and 201.1 mm at the maximum radial distance (i.e. the image corners), a difference of 0.45%. The 3σ uncertainty is ± 0.251 mm. Subsystem-level measurements showed significant geometric distortion in the WAC corresponding to 3.61 pixels in the image corners.

Observations taken during cruise after Jupiter encounter confirm the ground-based measurements of focal length and geometric distortion. In-flight observations using stellar targets (in this case, the Pleiades and open clusters M35 and M48) can constrain these values to much greater accuracy than in the lab. The WAC focal length was found to be 200.77 mm for CL1/CL2, and to vary slightly with filter combination, and the maximum distortion in the image corners was revised to 3.36 pixels. For the NAC, the measured focal length for CL1/CL2 was 2003.44 ± 0.03 mm, and the maximum geometric distortion was 0.45 pixels in the image corners. The same analysis found the WAC optical center at sample 548 and line 508, and the NAC optical center at sample 560 and line

500, albeit with a large 1-sigma uncertainty in the latter of about 30 pixels. However, field distortion in the NAC is essentially negligible, rendering this uncertainty largely moot.

The observed image pixel location (X_{im} , Y_{im}) can be related to its original, undistorted location (X_{corr} , Y_{corr}) as follows:

$$\begin{aligned} X_{corr} &= X_{im} * \frac{f'}{f} (1 + k * R_{opt}^2) \\ Y_{corr} &= Y_{im} * \frac{f'}{f} (1 + k * R_{opt}^2) \end{aligned} \quad (4)$$

where k is the camera distortion coefficient, f is the filter-dependent focal length, f' is the default or average focal length for that camera (assumed in the calculation of k), R_{opt} is the radius of the pixel from the optical center, and all distances are physical distances in mm. Focal lengths for NAC and WAC filters derived from the in-flight analysis described above have been provided in `nac_focallengths.tab` and `wac_focallengths.tab` files in the `calib/distortion/` subdirectory of the Calibration archive volume. Current best estimates for the distortion coefficient k are $-6.2 \times 10^{-5} \pm 3 \times 10^{-6}$ for the WAC (with a corresponding f' of 200.956 mm) and $-8 \times 10^{-6} \pm 2 \times 10^{-6}$ for the NAC (f' of 2002.61). To get a sense of how small this is, if instead $k = 0$ is assumed for the NAC along with an f' of 2003.181 mm, the maximum displacement in the image corners (assuming optical alignment) is 0.17 pixels. Geometric correction for the NAC is thus only recommended if extremely high sub-pixel accuracy is required.

The NAC optical system was also measured for its susceptibility to stray light – the fraction of light from an off-axis target that makes it into the detector when it is outside the field of view compared to the amount measured when the target falls within. The amount of stray light was found to be 3 to 4 orders of magnitude less than the maximum given in the flight specification, in both visible and UV wavelengths.

Dedicated measurements of stray light in both cameras were performed during the cruise portion of the mission and are described in West et. al. (2010). The pattern of stray light across the field

of view is caused by off-axis light reflecting off of surfaces in or around the cameras, and is a complicated function of the size and shape of the stray light source, its angular distance from the camera axis, and angular orientation about the axis. It is therefore quite difficult to model or remove in software. Stray light in the WAC is worse than in the NAC by about three orders of magnitude due to the WAC's refractive optical design, and is somewhat worse in the infrared than in visible wavelengths. Solar stray light is particularly problematic, and reaches an I/F of about 0.1 when the WAC is pointed 20 degrees from the sun; the detector saturated in test observations when pointed within 15 degrees. In practical terms, this means that the WAC is unusable for imaging most faint or diffuse targets at phase angles greater than 150 degrees.

Verification of opto/mechanical alignment of the boresight with the camera mounting surface was performed for both cameras during pre-flight calibration. All alignment parameters were found to fall within the required tolerances. Subsequent tests on the assembled spacecraft measured the alignment of the NAC boresight with all of the other instruments on the remote sensing pallet, as well as the two stellar reference units (SRUs/"star trackers"); these were also found to comply with specifications. This process is described in Section 5.3.2 of the Calibration Report.

Filters

Spectral transmission for all of the ISS flight filters was measured by the manufacturer, Barr Associates, as well as by JPL. The JPL measurements are described in Section 4.2 of the Calibration Report (and archived in Appendix F of that document), and are the ones used by CISSCAL for calibration. Two separate data sets are available: the "transmission" data is a set of measurements taken over the spectral region of interest at five separate locations on each filter, and then averaged; and an additional set of higher-resolution "blocking" data, measured at just one location on the filter (the center), but covering the entire spectral range of the instrument, from 200 to 1100 nm. The spectral

resolution is 1 nm in most cases. The system transmission files used by CISSCAL (contained in the calib/efficiency/systrans/ subdirectory of the COISS_0011 Calibration volume) were constructed by first combining the “transmission” and “blocking” data sets. The measurements were averaged in the overlap regions, and the blocking data were used to fill in the regions of low spectral response, before finally convolving the filter transmissions with the optics and detector quantum efficiency.

Uncertainty in the transmission measurements is thought to be about 1% over most of the bandpass, however, in regions where the signal falls below around 10^{-4} , it is much higher, as high as 50%. In addition, some particularly low-response regions where the filter transmission fraction falls below about 10^{-6} , approaching the sensitivity limit of the laboratory spectrophotometer, exhibit erroneous features such as abrupt jumps or negative values. These features occur at such low levels that they should not have a significant effect on the calculated system transmission.

A separate set of tests was required to measure the transmission of the polarized filters. This process was complicated by the fact that the light beam of the spectrophotometer is highly polarized. It was therefore necessary to use a reference polarizer in conjunction with the polarized filter being measured, and average together measurements taken at several orientations relative to the incoming beam in order to approximate a non-polarized light source. This process is described in Section 4.2.3 of the Calibration Report, and the resulting transmission measurements for the parallel (maximum transmission) and perpendicular (minimum) cases is provided along with the rest of the filter transmission data in Appendix F of that document. (Note that the polarizer transmission currently provided in the calib/efficiency/systrans/ subdirectory are an average of the parallel and perpendicular cases, and thus not particularly useful. Future versions of the Calibration archive volume will contain both orientations separately.) The polarizer orientations were also measured during subsystem-level calibrations, and the results of those measurements are provided in Table 6.

In-flight “red leak” measurements also help to constrain our understanding of the filter

transmission behavior. As described in West et. al. 2010, we can use long exposures of Vega taken with “canceling” filters – e.g. a UV filter in one filter wheel and an IR filter in the other - to measure the opacity of the combination and make sure it is commensurate with our knowledge of the individual filter transmissions. Our analysis found the signal in these images to be within about 25% of what we would expect given the laboratory-measured filter transmissions, which is consistent with the known uncertainties.

Filter Combination	Angle of Max. Response	Percent Error
P0 / CL2 (NAC)	-0.5	0.566
P60 / CL2 (NAC)	61.8	0.580
P120 / CL2 (NAC)	120.8	0.434
IRP0 / CL2 (NAC)	2.25	1.3 4
CL1 / IRP0 (WAC)	0.00	0.22
CL1 / IRP90 (WAC)	90.85	0.18

Table 6: ISS polarizer orientations as measured in the laboratory.

Radiometric flatfields were acquired during thermal vacuum testing for each filter combination. The goal is to compare the independently-measured radiance of an evenly-illuminated light source to the measured DN level across the detector. After accounting for the effects of shutter offset, dark current/RBI and bias level, the slope term representing the best linear fit to the detector response is obtained. These are recorded in VICAR-format “slope files” in the calib/slope/ subdirectory of the COISS_0011 volume. The actual values recorded are the reciprocal of the slope term, in units of picoamp-milliseconds/DN. (Here, picoamps refers to the current measured by the light meter, and can be converted to radiance units using the data tables supplied in Appendices F13 and F15 of the Calibration Report.) The images in this analysis were acquired with anti-blooming mode turned on; the slope files used by CISSCAL have had anti-blooming pixel pairs removed. Flatfield images corrected

with the derived slope files are typically flat to within a few percent. Residual artifacts caused by internal reflections remain visible at or below the 1% level.

Of particular note are regions of low sensitivity observed in the flatfield of the WAC VIO filter. These are likely due to imperfections in the Lumogen coating added to the CCD to boost UV sensitivity. This pattern of low sensitivity is not seen in any other filter combination.

As with many of the ground calibration measurements, the light transfer measurements were taken at three different optics temperatures, +5, -10, and +25 degrees Celsius, representing the nominal value, lower operating threshold, and upper operating threshold of the instrument respectively. These measurements showed significant variation in flatfields taken at different temperatures, the cause for which remains unknown. Since the vast majority of ISS data is taken near the nominal operating temperature, the +5 degree flatfields will be used for calibration in almost all cases.

The `in_flight.pdf` document included in the Calibration volume document/ directory details the strategy for in-flight measurements of the flatfield. In theory, these are obtainable during close flybys of relatively featureless targets like Titan; however, in practice, the large number of images required to average over low-frequency spatial variations, combined with the large number of filter combinations required – 81 in all – make this difficult to achieve. In-flight analysis has been performed to identify flatfield changes from launch through rev 225; initial efforts are documented in West et. al. (2010). One dust ring seen in flight images in all NAC filter combinations has been characterized using data from Venus flyby. Additionally, a “mottling” correction for the NAC has been derived from Titan images, and accounts for low-level spatial variations seen after SCLK=1444733393. These corrections are applied by CISSCAL using data files contained in the `calib/dustring/` subdirectory.

Shutter

ISS exposure times are controlled by the shutter mechanism as described in Section 2.3 of this

User's Guide. The actual exposure time differs from the commanded time by a constant offset as well as a variable offset across the field of view. The latter is referred to as “shutter shading” and is caused by the movement of the shutter blades in the sample direction. Component-level measurements of the shutter mechanism were used to configure the instrument at the subassembly level, by adding a constant offset to the commanded exposure time in software, and by positioning the shutter mechanism such that the CCD window falls on the smoothly-varying portion of the shutter shading plot. The sample-dependent offset was then measured in thermal vacuum testing. The resulting shutter offsets for temperatures of -10, +5, and +25 degrees Celsius have been recorded to VICAR-format image files and included in the `calib/offset/` directory of the Calibration volume.

In-flight observations of stars have been used to update the constant shutter offset value, which has changed since launch. The analysis consists of imaging a target of constant brightness at a range of exposure times from 5 ms (the minimum non-zero exposure) up to as long as is possible without saturating. This is a bit tricky to accomplish in a single filter combination due to the limited dynamic range of the instrument. Current best estimates using full-disk images of Rhea peg their constant offset values at 2.75 and 2.67 ms respectively; that is, the actual exposure times for the NAC are 2.75 ms less than commanded, and those for the WAC are 2.67 ms less than commanded. These offsets are hard-coded into CISSCAL's exposure correction algorithm and adjusted as new data becomes available; note that these corrections are applied only in calibration and not to image timing keywords such as `IMAGE_MID_TIME`. The latest analysis indicates a very small intrinsic shutter time uncertainty in both cameras of about 0.03 ms, which is less than 1% of the shortest commandable exposure. However, there is currently some question as to whether, and to what extent, the shutter offset actually drifts over time; this will require more investigation to resolve.

Detector

The CCD detectors used by the ISS are front side-illuminated, 1024x1024 pixel arrays, with pixels 12 μm to a side. The CCD chips were manufactured by Loral and packaged by JPL in a hermetically sealed enclosure behind a fused silica window. Light leakage was minimal to nonexistent for the equivalent of 1000 seconds of exposure on Saturn. The CCDs have a Lumogen coating to enhance UV response. The operating temperature of the detectors in flight is -90°C , and is maintained via active and passive heating elements. There are also decontamination heaters for removing outgassed volatiles that can accumulate on the CCD window. (See the description of the NAC haze anomaly in Section 3.6 of this document.)

Quantum efficiency (QE) is defined as the number of electrons generated per incident photon. Expressed as a function of wavelength, it describes the spectral sensitivity of the CCD detector. QE measurements were performed in two independent tests, one which was considered more reliable but limited in wavelength coverage, and another which was less reliable but spanned nearly the entire spectral response. These two data sets were combined, using a weighted average in the overlap region, into the data files provided in Appendix F of the Calibration Report, and the `calib/efficiency/ccdqe.tab` file used by CISSCAL. As can be seen in Figure A.6 in the Appendix of this document, the QE of both cameras plateaus between about 550 and 750 nm at ~ 0.35 electrons per incident photon. Quantum efficiency is difficult to measure accurately, and so this quantity is thought to be the most poorly-constrained of the variables comprising the system transmission, with uncertainties as high as 20%.

For an ideal detector, the square of the noise measured in a particular region plotted against the mean DN value will exhibit a linear relationship according to:

$$\sigma_S^2 = \frac{(S - R)}{g} + \sigma_R^2 \quad (5)$$

where S is the mean signal measured in DN, R is the mean bias level, g is the system gain, σ_S^2 is the

variance in S , and σ_R is the read noise floor. When σ_S^2 is plotted against the mean signal (corrected for bias) the slope of the line thus represents the inverse of the system gain, and the turnoff point at which the DN values start to deviate from linearity provides a rough estimate of the full-well capacity.

Variations on this test were performed during both the component and subsystem-level testing, and are described in Sections 4.4.3 and 5.1.2 of the Calibration Report. A least-squares fit to the equation above gives a baseline value for gain state 2 of 30.27 e⁻/DN in the NAC and 27.68 e⁻/DN in the WAC. (The latter is an average of the anti-blooming mode ON and OFF cases.) Best estimates indicate a full-well capacity for un-summed images of close to 110,000 electrons in the NAC, and between 90,000 and 100,000 electrons in the WAC. For gain 2, these correspond to signal levels of around 3600 and 3400 DN, respectively, well under the 12-bit saturation level of 4095 DNs. In addition, measurements taken in 4x4 summation mode suggest that the summation well – a separate entity from the pixel wells and not simply the sum of the constituent full-well capacities – starts to saturate around 1000 DN.

The same data set used to estimate the full-well capacity and gain at the subsystem level was also used to measure gain ratios, sensitivity, linearity and non-linear response terms, and to identify “blemish” pixels. This data set (OBSERVATION_ID = 'LTC_BLEM_GAIN') consists of images taken at increasing and decreasing exposure times at each of the standard summation/gain state pairings, with illumination provided by a flat target with known radiance. Images were acquired under thermal vacuum at temperatures of +5 and +25 degrees Celsius.

Gain ratios were determined by measuring the system sensitivity, i.e. the slope term in the linear equation relating the signal in DN to the incident energy. Energy is computed as the measured radiance multiplied by the exposure time. Test results showed no significant dependence of sensitivity on either temperature or anti-blooming mode. The ratio of the fitted sensitivity slope values for different gain states, averaged over 100 separate areas of the image, provide the nominal gain ratios. Combining these

values with the baseline gain 2 values derived from the “noise squared versus signal” analysis described above yields gain values for each mode.

The camera gain ratios have since been further refined by Bob West using Saturn images from cruise by comparing images with different gain states taken close together in time. The values derived from this analysis are provided in Table A.4 in Appendix A of this document and provide the current best estimates of gain for gain states 0, 1 and 3.

The linearity of the instrument response to light with respect to the best-fit sensitivity values is described in Section 5.1.11 of the Calibration Report. Non-linear correction terms were derived to convert the measured DN values to their theoretical values assuming a linear DN response with exposure time. Since the behavior at low DN is expected to be most linear, a weighting scheme was chosen to give additional weight to this region. The results have been incorporated into the linearity correction module of CISSCAL. Significant non-linearity is seen in 4x4 summation mode/gain state 0 above a DN level of 1000. Additional analysis of these data (Section 5.1.3) identified “blemishes,” pixels exhibiting highly non-linear behavior beyond specified thresholds. The vast majority of these blemishes were found to be located along the edges of the CCD array, in either the first or last line or sample.

As discussed above, setting anti-blooming mode to ON produces an effect in long-exposure images whereby electrons accumulate in traps in certain pixels at the expense of charge in the pixel adjacent in the negative line direction. These bright/dark pixel pairs are described in Section 5.1.10 of the Calibration Report. VICAR-format maps were made for the pixel pairs that behaved more-or-less consistently – this can be found in the calib/antibloom/ directory of the COISS_0011 Calibration volume. These maps proved to be inadequate for removal of the pixel pair effect in flight images, so CISSCAL currently employs a method to identify pairs by using a simple threshold detection algorithm on the image itself. The algorithm replaces the erroneous pixel values with the mean of their adjacent

neighbors in the sample direction.

The ISS detectors also suffer from an effect known as uneven bit weighting. This effect is described in Section 5.1.9 of the Calibration Report as follows:

“Observing the histogram of an arbitrary image with a broad range of DN shows a periodic series of spikes centered at $(2^m) \cdot n$ (where m and n are integers) and occurring over a span of several DN. The largest of the spikes appear where m , the bit level, is large, or at the greater significant bits in the digitized output. The cause for this is attributed to the comparator in the analog to digital (A/D) converter. The A/D converter used in both ISS cameras [...] uses the process of successive approximations to digitize the data[...] During this process, the specific DN values at 2^m (e.g. 2048) are digitized by comparing the reference voltage across a single resistor, whereas nearby DNs (e.g. 2047) are digitized with the adjustable resistor network appropriately set. Switching between the two DNs seamlessly requires balancing the group of resistors to the single resistor to an accuracy of 1 part in 10^4 , which is difficult to achieve in a practical device. Thus the imbalance in resistances produces the spikes.”

This uneven distribution of DN values can be modeled and corrected mathematically. This analysis employed a device called a Dynamic Ramp Target to create a smoothly-varying distribution of DN values that is constant with sample number, across the entire dynamic range. A “super histogram” is then constructed by combining the histogram distributions of many individual images. Normalizing this histogram to an ideal histogram derived from a filtered version of the original produces an “effective bin width” for each DN value. Adding these bin widths sequentially gives an array of adjusted DN values that can then be applied in software to the observed DNs as the very first step in the calibration pipeline (after conversion from 8 to 12-bit, if necessary). Note that this correction requires converting from an integer to a floating-point format to accommodate fractional DN values.

Files containing the adjusted DN values for the appropriate combination of gain state and

temperature are located in the calib/bitweight/ subdirectory of the COISS_0011 calibration volume.

Data from 0 to 200 DNs is left uncorrected due to problems with light leaks corrupting the data in that region. Also note that the bit weighting analysis was cut short during the testing process, and that with the exception of the 1x1/gain 2 case, all NAC data was acquired at a temperature of -10° C rather than the nominal +5° C. This may adversely affect the quality of the bit weight correction for those gain states, as significant variability with temperature was observed.

Section 2.1 of the ground Calibration Report describes a “dark band” running along the left side of NAC images, affecting samples 1 through 12. The WAC is unaffected. Laboratory investigations traced the problem to abnormally high resistance between the CCD and ground caused by imperfect bonding of the CCD chip to the ceramic header. This problem also affects the Extended Pixels and DARK_STRIP_MEAN in the NAC, rendering them unusable.

The noise properties of the detectors are discussed in Section 5.1.6. Fourier analysis reveals low-level fixed pattern noise in both the horizontal and vertical directions, and out-of-phase electrical noise in the horizontal direction. In addition, setting the anti-blooming mode to ON was found to boost the amplitude of vertical spatial frequencies corresponding to 0.1 Hz and above. This effect becomes more noticeable as the signal level increases.

After launch, both cameras began picking up a more significant low-frequency fixed-pattern noise source consisting of two main peaks in the 2-3 Hz range in the NAC, and one main peak near 4 Hz in the WAC. The exact origin of the noise is unknown, but it seems to be introduced by the analog-to-digital converter (ADC) during readout and is presumed to originate in some electrical component of the spacecraft. The pattern varies by only a few DN (in gain state 2) and so is most noticeable in images with intrinsically low signal levels, e.g. dark sky and dark frames taken with the shutter disabled (see Figure 8). In the two lowest gain states, often used in conjunction with 2x2 and 4x4 summation, the 2 Hz noise is present only at a very low level - about 1 DN in gain state 1, and 0.5 DN

in gain state 0. Also, in summed images, the noise manifests as a more complicated diagonal pattern rather than a coherent horizontal pattern.

A filtering algorithm has been incorporated into CISSCAL to remove this noise pattern in unsummed images, using either the overclocked pixels or, ideally when possible, a horizontal mean of dark sky image values. More information about the 2 Hz removal algorithm and its implementation can be found in the `theoretical_basis.pdf` as well as Section 4.4 of this document, below.




Figure 8: Example of the 2 Hz horizontal banding pattern in a NAC dark image. This particular image was too large to fit completely in the image buffer during readout, resulting in a significant change in the readout rate half-way through. From Porco et. al. (2004).

As discussed above, the ISS detectors exhibit very low dark current at their nominal operating temperature of -90°C , but significant residual bulk image (RBI) effects caused by “trapped” electrons from the previous exposure leaking out into the pixel well during the subsequent exposure. Section 4.4.4 of the Calibration Report and Porco et. al. (2004) describe how the contribution of RBI to a given pixel can be modeled as a sum of exponentials:

$$DN_{RBI} = a_1(1 - e^{-b_1t}) + a_2(1 - e^{-b_2t}) + a_3(1 - e^{-b_3t}) \quad (6)$$

Here, each exponential term represents a “trap” caused by a bulk-state defect within the silicon. The coefficients a and b represent the trap size and time constant, respectively. A sum of three traps was deemed sufficient to model the RBI accurately. Adding in constant and linear terms to represent the standard dark current contribution yields the following 8-parameter model:

$$DN_D = a_1(1 - e^{-b_1t}) + a_2(1 - e^{-b_2t}) + a_3(1 - e^{-b_3t}) + ct + DN_0 \quad (7)$$

which can be solved for each pixel via least-squares, given enough data taken over a large enough range in exposure times.

The current dark current implementation in CISSCAL, described by West et. al. (2010) uses a more direct method than the one described above. Rather than fitting to a model, we can simply construct a measure of the average dark current/RBI over a range of exposure times and then interpolate in the time domain. This process is complicated by the fact that the dark current/RBI accumulates in a pixel not only during exposure, but during readout as well: as the charge from a particular potential well is shifted vertically, one line at a time, towards the readout register, it picks up additional charge from each physical pixel location through which it passes. Thus, we can express the RBI contribution for a particular pixel originating at position $[i,j]$ as the sum of the contributions from that pixel and every pixel “downstream”:

$$RBI(i, j) = \sum_{k=0}^j RBI(i, k, t_2) - RBI(i, k, t_1) \quad (8)$$

where t_1 and t_2 correspond to the times at which the potential well enters and leaves physical location $[i,k]$. Precise calculation of these times for each line on the CCD is a complicated function of many camera parameters, and is accomplished by a specialized “linetime” code, included with the CISSCAL

distribution. The effect of readout on the dark/RBI pattern is especially apparent in 1x1 uncompressed and lossless-compressed 12-bit images, which contain more bits of information than can fit in the readout buffer at once, resulting in an abrupt change in readout rate half-way through image readout, and thus a discontinuous increase in the slope of the dark/RBI signal with line number.

Data files containing average dark counts for each pixel measured at 8 discrete exposure times (0, 10, 32, 100, 220, 320, 460, and 1200 seconds) are contained in the `calib/darkcurrent/` directory of the Calibration volume, and provide the raw data for the interpolation. Dark calibration observations are planned at intervals throughout the Cassini mission, and the parameter files are named according to the time (year and day-of-year) of the observation from which they are derived. Because the dark simulation is computationally intensive, simulated dark files for a given set of camera parameters are saved for future use (to subdirectories of the `calib/darkcurrent/` directory, named according to the source dark parameter file). CISSCAL first checks these directories for an appropriate simulated dark file before creating a new one. Simulated darks are stored in VICAR image format.

In general, the dark/RBI pattern is seen to be consistent over time, however, some pixels exhibiting abnormally high electron emission rates have been found to appear, or in some cases disappear, from one dark calibration observation to the next. We call these “hot pixels.” The exact mechanism by which this occurs is unknown, but the working theory suggests that the hot pixels are simply caused by particularly large defects in the chip – perhaps created by collisions with cosmic rays or some other high-energy particle - which can later sometimes be “healed” through annealing of the silicon. CISSCAL 3.7 incorporates code that corrects for changes in the hot pixel distribution over the course of the mission.

Charge transfer efficiency (CTE) is a measure of the fraction of charge retained by a pixel during the readout process. Pre-flight measurements calculated a CTE for both ISS detectors of 0.99994. This value is expected to decrease somewhat over the mission, and so in-flight observations

have been designed to look for any changes. As described by West et. al. (2010) these involve imaging star clusters (M48 for the NAC, Pleiades for the WAC) at 4 orientations offset 90 degrees from one another, and plotting any change in star flux as a function of the difference in line value. Unfortunately, the measurement error in the star photometry is such that attempts to derive CTE using this technique have so far been unsuccessful.

Point-Spread Function

The point spread function (PSF) or point response function (PRF) is the cross-sectional detector response generated by a point source of light passed through the optical system. It is a function of both camera and filter combination. The PSF width, generally reported in terms of the full-width at half-maximum (FWHM), provides the fundamental measure of the resolving power of the instrument. In the ISS, PSF widths are typically 1-2 pixels; in other words, the PSF is not fully resolved by the pixel sampling size. Laboratory measurements of the ISS PSF conducted during pre-launch testing had insufficient dynamic range to measure the point response out to the edge of the CCD. In-flight measurements are preferable in any case, due to the possibility of PSF changes during launch or cruise, or caused by events such as the 2001 NAC haze anomaly (described in Section 3.6 of this document).

In-flight measurements of the PSF are conducted by averaging together many short-exposure images of bright stars. The reason for the averaging is twofold: 1) to suppress the background noise and thus boost the dynamic range, and 2) to allow sampling of the PSF at a sub-pixel resolution, by moving the star on the detector between exposures, and interpolating to a finer grid before aligning and combining. Dynamic range can be extended by using saturated stars to characterize the outer PSF region, and combining this with a core PSF derived from well-exposed images. West et. al. (2010) describes how further enhancement of the extended region can be achieved using satellite images, through a process of iterative convolution. A summary of these in-flight measurements – the FWHM

and dynamic range for each measured filter combination - is given in Table A.2 in the Appendix, and an example of a typical PSF is shown above in Figure 9. At the time of this writing, the PSFs generated by this analysis have yet to be archived, but are planned for inclusion with future releases of the COISS_0011 volume.




Figure 9: Point-spread function for the NAC BL1/CL2 filter combination, shown here on a base-10 logarithmic scale. From West et. al. (2010).

A handful of filters, including the NAC GRN filter, exhibit low-level “ghost” peaks, or other asymmetries in their PSF (see Figure 10). Also notable is the WAC IR1/IR2 combination, which exhibits an abnormally wide PSF due to a quirk of the instrument design, whereby focus is optimized for filters when they are used alongside one or the other clear filter.




Figure 10: Point-spread function for the NAC BL1/GRN filter combination, plotted on a logarithmic scale. Note the faint off-axis secondary peak. From West et. al. (2010).

Compression

The details of the compression algorithms used by the ISS lie beyond the scope of this manual. Interested readers should consult Sections 5.5.1 and 5.5.2 of the Calibration Report, which provide full descriptions of the lossy and lossless compression schemes used by ISS, and report on pre-flight testing to verify the compression efficiency versus image entropy, and to measure the amount of data loss in the lossy case.

Data truncation caused by compression-related effects is discussed in Section 3.6 of this document. Note that use of the lossy compression scheme has been discontinued in the XXM in order to conserve downlink processing work-hours.

Absolute Calibration

Finally, we arrive at the important subject of absolute flux calibration, the final step in the radiometric calibration sequence described in Section 4.2 above. This discussion will focus on how to derive the absolute radiometric correction factors for each filter combination, which is to say, the C_{f_1, f_2} term in Equation 2. It is possible to derive a set of correction factors from the pre-flight flatfields, which have been calibrated for absolute brightness (the “slope” files are so-called because they contain the slope term describing the linear relationship between DN and absolute radiance), but in-flight measurements of photometric standards are preferable, since absolute sensitivity can change over time due to accumulation of outgassed volatiles on the optical elements or other unknown causes. To monitor these potential changes, absolute calibration observations are scheduled on a roughly annual (in Earth-time) basis throughout the duration of the Cassini mission.

Stars with well-described spectral profiles are used to establish the absolute flux scale; the Vega spectrum from Bohlin & Gilliland (2004) is the primary stellar reference for ISS. Fainter stars are also imaged in order to allow longer exposures for broadband filters and thus minimize the contribution from shutter uncertainty; as well as hot O and B-type stars, for their higher flux in the UV (since A and G stars exhibit significant falloff at wavelengths shortward of the Balmer discontinuity at 365 nm). Stellar variability, uncertainty in the reference spectra, and possible sub-pixel effects all contribute error to the photometry results. In addition, the spectral shape of stars can be quite different than that of planetary targets, making them a less than ideal standard of comparison. Luckily, we can constrain the star analysis results using additional methods: satellites, especially Enceladus, provide a bright, flat spectral response which can be used to help establish relative color balance, and a comparison of simultaneous (BOTSIM) images in equivalent filter combinations provide a way to calibrate the cameras with respect to one another.

The most recent radiometric correction analysis using in-flight data collected through 2009 is

described by West et. al. (2010), and will be explained in great detail in the end-to-end example in Section 5; here we will briefly summarize the procedure. All images were processed in CISSCAL according to the steps laid out in Section 4.2, with default settings used for most calibration options. The notable exception is the 2 Hz noise removal, for which the image mean method was used, as the 2 Hz pattern shows up well against images with a dark sky background. Bias was subtracted using the BIAS_STRIP_MEAN value, as is generally recommended in conjunction with image-mean 2 Hz subtraction.

This analysis is simplified by built-in functionality in CISSCAL that can normalize the output flux to an input spectrum specified by the user (the “F” in I/F). A median boxcar filter is used to remove cosmic ray spikes. For stars, aperture photometry is then performed on these spectrum-normalized, calibrated images; if the instrument parameters are well-known, the total integrated flux from the star should sum to a value near unity that represents the ratio of the measured flux to the reference spectrum flux. Averaging the result for several images taken with the same filter settings provides our measure of $C_{f1,f2}$.

Satellite images are treated similarly to stars. We first calculate the global mean albedo by calibrating the images to I/F, adding up all the light from the satellite and dividing by the number of pixels subtended by the disk. We then normalize this to the known albedo spectrum – in this case an HST spectrum of Enceladus provided by Keith Noll (2008, private communication) – which is convolved with each filter combination bandpass to produce comparison (theoretical) filter-integrated global albedo values. Modulo corrections for phase and surface brightness variations, the ratio of the measured global mean albedo to the comparison albedo provides an alternate measure of $C_{f1,f2}$, which we use to constrain the relative colors. Unfortunately we are not able to use Enceladus to constrain the absolute flux scale, as our measured Enceladus fluxes are systematically low for all filters by about 20%. This discrepancy is so far unexplained.




Figure 11: Corrected Vega flux for NAC and WAC as a function of filter effective wavelength. From West et. al. (2010).

The quantum efficiency measurements taken in the lab are thought to be the most uncertain of the contributions to system transmission in Equation 2. This suggests we can derive a first-order correction to the quantum efficiency by plotting $C_{f1,f2}$ versus filter effective wavelength, interpolating to a fine wavelength grid, and smoothing to minimize the high-frequency filter-to-filter variations. The resulting QE correction function for each camera, Q_{corr} , is stored in data files (`nac_qe_correction.tab` and `wac_qe_correction.tab` in the `calib/efficiency/` subdirectory of the `COISS_0011` volume) which are read by `CISSCAL` and applied to the ground-based quantum efficiency during calibration. The residual offsets then become the new $C_{f1,f2}$ factors. For the West et. al. (2010) analysis (see end-to-end example in Section 5), this process was repeated a second time in a slightly different fashion: residual correction factors for both WAC and NAC were plotted against effective wavelength and a line was fit to the

entire set of data as a whole. This slope was then divided out of the quantum efficiency corrections for both cameras, ensuring that the camera-to-camera ratios between identical filters remained unchanged. New, final residual $C_{f1,f2}$ values were then computed and the results saved to the `correctionfactors_qecorr.tab` file in the `calib/correction/` subdirectory.

Tables 2 and 3 in West et. al. (2010) describe the star targets and satellite data, respectively, used to derive the absolute correction factors for CISSCAL 3.6, and Figure 11 shows corrected Vega fluxes for the NAC and WAC overlaid with the reference spectrum from Bohlin & Gilliland (2004). Error estimates are provided in Tables 4 and 5. These errors were calculated empirically by re-running the photometry analysis on reference spectrum-normalized, absolute flux-corrected images and calculating the standard deviation for each filter. Note that here the Enceladus data was scaled to the average Vega brightness, which in our normalized frame of reference is unity.

Figure 12 contains a plot of normalized Vega flux for all filters over the course of the mission. Images with exposure times less than 25 ms have been excluded to reduce the effect of shutter uncertainty. No obvious long-term trends are observed on this timescale.

Extensive detail about the derivation of the radiometric correction factors in CISSCAL can be found in the end-to-end example in Section 5. This discussion includes image search strategies, a detailed description of the star and satellite photometry, methods for excluding erroneous data and combining results, and other procedural details.

A final note about absolute calibration: a separate photometric analysis has been performed using full-disk and 2x2-mosaiced images of Jupiter from Cassini's Jupiter flyby. The results of this analysis, which uses the 300-1000nm spectra from Eric Karkoschka (1994, 1998) are different enough from the analysis described above that we recommend ISS Jupiter images be calibrated using a separate set of absolute correction factors. These correction factors are provided in the `calib/correction/jupiter_correction.txt` file included with CISSCAL 3.7, and a separate calibration

option has been incorporated into the absolute calibration step in the software interface. This alternate correction applies to the BL1, GRN, CB1, CB2, CB3, MT1, MT2, and MT3 filters in the NAC (paired with clear) and the VIO, GRN, CB2, CB3 and RED filters in the WAC. The exact cause of this discrepancy is still under investigation.




Figure 12: Corrected normalized intensity of Vega for all filters with exposures greater than 25 ms, plotted against time over the course of the mission. Lines have been fit to the average values for each filter. The absolute calibration as a whole appears stable on this timescale.

Polarimetric Calibration

The use of the polarizing filters to obtain polarimetric quantities is described in West et. al. (2010). The NAC contains three polarizers with orientation offsets of 60 degrees. They can be used together to determine intensity, degree and angle of linear polarization. The WAC contains two orthogonal polarizers that can be used to determine intensity and the Stokes parameter Q. Included with

CISSCAL 3.7 is a standalone code called `make_polar_image.pro` which can be used for calculating these quantities according to the methods discussed in this section; see also Section 4.4 for more information on usage of this code.

Cassini ISS visible polarizers are made of polaroid material; IR polarizers are wire grid. In both cases the transmission through the filter can be described by Equation 9:

$$I_t = \frac{I_u(T_{\parallel} + T_{\perp})}{2} + I_p(T_{\parallel}\cos^2(\vartheta) + T_{\perp}\sin^2(\vartheta)) \quad (9)$$

where T_{\parallel} and T_{\perp} are the polarizer transmissions parallel and perpendicular to the principal polarization direction, I_u and I_p are the unpolarized and linearly polarized components of the incident intensity and ϑ is the angle of the polarized electric vector with respect to the principal polarization direction of the polarizer.

Prior to assembly of the filter wheel T_{\parallel} and T_{\perp} were measured in the laboratory, but only for one filter.

We have measurements of Titan taken on day 90 of 2009. The spacecraft was rotated into five positions such that the sun azimuth was -90, -120, -150, -180 and 150 degrees. From Equation 9, plus measurements of I_t using a clear filter in place of the polarizers, we can solve for the sum of the parallel and perpendicular transmissions for each polarizer. First we recast Equation 9 into the following form:

$$I_t = \frac{I_u(T_{\parallel} + T_{\perp})}{2} + I_p(T_{\parallel}\cos^2(\vartheta) + T_{\perp}(1 - \cos^2(\vartheta))) \quad (10)$$

We can then perform a least-square linear fit of the measurements (polarizer DN values divided by clear-filter DN values) to Equation 10 with $x = \cos^2(\vartheta)$ as the independent variable. The slope gives $I_p/I_t * (T_{\parallel} - T_{\perp})$ and we can obtain $T_{\parallel} + T_{\perp}$ from the fitted values at $x = 0$ and 1.

If we were able to measure $T_{\parallel} - T_{\perp}$ instead of $I_p/I_t * (T_{\parallel} - T_{\perp})$ we could derive T_{\parallel} and T_{\perp} . If we assume, as laboratory measurements show, that for the IR polarizers $T_{\perp} \cong 0$, we are then able to derive T_{\parallel} . Analysis of the data shows that if $T_{\perp} = 0$ for both polarizers the two IR polarizers give different values for I_p . Therefore a value of 0 was assigned to only one of the polarizers and the value of T_{\perp} for the other IR polarizer was calculated to make I_p the same for both polarizers. These steps lead to values for parallel and perpendicular transmissions for both polarizers. Knowing these values Equation 9 can be solved for both polarizers to give:

$$I_u = \frac{2(I_{90}T_{\parallel,90} - I_0T_{\perp})}{D} \quad (11)$$

$$Q = \frac{I_0(T_{\parallel,90} + T_{\perp,90}) - I_{90}(T_{\parallel,0} + T_{\perp,0})}{D} \quad (12)$$

where

$$D = T_{\parallel,0}T_{\parallel,90} - T_{\perp,0}T_{\perp,90}$$

and I_0 and I_{90} are the measured intensities in the polarizers at 0° and 90° orientations. In Equation 12 Q is the Stokes parameter Q defined in the reference frame of the polarizers. In practice we express Q as a ratio of Q defined above to the total intensity, I_t (the sum of all terms in Equations 11 and 12). With that definition Q has the range $-1 \leq Q \leq 1$ and is the same as the polarization if the electric vector is aligned with the polarization axis of one of the polarizers (as is often the case for polarization observations). Table 7 gives the values of the transmissions for the IR polarizers derived from the Titan observations.

It is also possible to obtain Q from images combining a color filter with CLR (clear) and IRP0 or P0. In that case I is obtained directly from the CLR filter combination and:

$$Q = I + \frac{2(IRP0 - I * T_{\parallel,0})}{(T_{\parallel,0} - T_{\perp,0})} \quad (13)$$

Analysis of the Titan data for the three visible polarizers also relied on linear fits to the signal as a function of $\text{Cos}^2(\vartheta)$. Figure 13 shows that no two polarizers have the same slope and intercept, although the fits are similar for the P60 and P120 polarizers. The P0 fit is, to first order, offset vertically from the other two and has a slightly different slope. This behavior, shown in Fig. 13 for the BL2 filter, holds true for the polarizers paired with the other passband filters as well.

Camera/Filter	$T_{\parallel,0}$	$T_{\perp,0}$	$T_{\parallel,90}$	$T_{\perp,90}$
WAC/MT2	0.828	0.01607	0.819	0.0
WAC/CB2	0.980	0.02279	0.905	0.0365
WAC/CB3	0.947	0.04165	0.967	0.0084

Table 7: Derived transmissions of ISS IR polarizers.

It is possible to derive an algebraic solution to invert measurements in the three polarizers to give I, P and ϑ if one can write $T_{\parallel,120} = T_{\parallel,60}$, $T_{\perp,120} = T_{\perp,60}$, $T_{\perp,0} = T_{\parallel,60} + \delta$ and $T_{\perp,0} = T_{\perp,60} + \delta$. Such a condition is plausible if, for example, a small chip or scratch removed some of the polarizing material in polarizer P0. Our retrieval algorithm does just that but the results provide a starting point for a more accurate iterative method. The iterative method uses the derivatives of Equation 10 to estimate how small changes in I_u , I_p and ϑ produce changes in each of the polarizer images. The result is a 3X3 matrix of sensitivity values that can be inverted to calculate what changes to I_u , I_p and ϑ are needed to produce agreement with observations in three polarizers. In practice the method is seen to converge to the desired accuracy after two iterations.




Figure 13: The ratio of the signal in each of three visible polarizers to the signal in the CLR filter (each paired with the BL2 filter) are shown as a function of $x = \text{Cos}^2(\vartheta)$. The measurements are shown as symbols for the three polarizers and the straight lines are linear fits to the measurements.

In order to use the iterative method one must have values for T_{\parallel} and T_{\perp} for each polarizer/passband filter combination. As noted above, the calibration measurements give the sum of these in each filter but only the product of the difference and the intrinsic polarization of Titan at the geometry of the observation. The intrinsic polarization of the target is same for each observation, so we can calculate P assuming one of the filters has $T_{\parallel} - T_{\perp}$ as measured in the lab, and then use that value of P to calculate $T_{\parallel} - T_{\perp}$ for the other two filters. That will ensure that all filters give the same value for P and that they can be use in combination, although the true value of P is still uncertain by the

amount that the true value of $T_{\parallel} - T_{\perp}$ used to do the normalization differs from the laboratory measurement. Values derived in this way are listed in Table 8.

Camera/Filter	$T_{\parallel,0}$	$T_{\perp,0}$	$T_{\parallel,60}$	$T_{\perp,60}$	$T_{\parallel,120}$	$T_{\perp,120}$
NAC/UV3	0.582	0.130	0.538	0.095	0.527	0.093
NAC/BL2	0.632	0.076	0.598	0.056	0.584	0.057
NAC/GRN	0.675	0.063	0.646	0.045	0.633	0.047
NAC/MT1	0.685	0.049	0.649	0.029	0.637	0.031
NAC/CB1	0.685	0.053	0.649	0.030	0.639	0.034
NAC/MT2	0.817	0.163	0.794	0.085	0.781	0.097
NAC/CB2	0.841	0.202	0.814	0.129	0.813	0.123

Table 8: Derived transmissions of ISS visible polarizers.

4.4 CISSCAL: The Cassini ISS Calibration Software Pipeline

Introduction

Note: The section that follows is an updated version of the `cisscal3_6_manual.pdf` document currently distributed with the `COISS_0011` Calibration archive volume, which this User's Guide replaces henceforth.

CISSCAL is the Cassini Imaging Science Subsystem Calibration software. It is essentially a graphical and command-line interface for performing the steps outlined in Section 4.2, using the data files described in Section 4.3. It performs standard CCD calibration steps such as bias and dark subtraction and flatfield correction, as well as ISS-specific calibrations such as bit-weight correction and removal of 2-Hz noise. CISSCAL only reads and writes images in VICAR format.

CISSCAL was developed in stages, beginning in August 1998, under the auspices of the Cassini

Imaging Central Laboratory for Operations (CICLOPS) directed by Dr. Carolyn Porco, ISS Team Leader. The theoretical basis and pipeline design for Cassini image calibration was developed by ISS team member Dr. Robert West of JPL. The systems design for the CISSCAL software, and its original implementation and development, was the work of Dr. Kevin Beurle at Queen Mary, University of London from August 1998 through March 2002. Beginning in April 2002, major design and algorithmic modifications, software additions, debugging and maintenance have been performed by Benjamin Knowles of CICLOPS/Space Science Institute, Boulder, CO under the direct supervision of West. Significant contributions (algorithms, software, evaluation and validation, etc) have been made by West, Vance Haemmerle of JPL, Daren Wilson of CICLOPS, and other members of the Cassini ISS Team.

All questions regarding CISSCAL should be directed to Ben Knowles at the Space Science Institute, Boulder, CO (ben.cisscal@gmail.com).

Setting up the Environment

CISSCAL is written in the Interactive Data Language (IDL), and thus requires that IDL (version 5.5 or later) be installed on the computer on which CISSCAL is to be run, and its executable placed in the user's PATH.

Once IDL has been installed, the user needs only to edit his or her login shell script (.cshrc, .tcshrc or similar) to define the CISSCAL-specific environment variables `CisscalDir`, `CalibrationBaseDir`, and `ImageBaseDir`. Assuming that the user is running the csh or tcsh shell and has installed CISSCAL in his or her home directory in a subdirectory named "cisscal3_8," and that the calibration support directory CALIB has been downloaded, the following lines should be added to the .cshrc/.tcshrc file:

```
setenv CisscalDir ~/cisscal3_7
```

```
setenv CalibrationBaseDir ~/CALIB
setenv ImageBaseDir ~/images
```

Specifically, `CisscalDir` specifies the location of the CISSCAL software files, `CalibrationBaseDir` specifies the location of the calibration support directories (version 3 of this software requires: `antibloom/`, `bitweight/`, `correction/`, `darkcurrent/`, `distortion/`, `duststring/`, `efficiency/`, `lut/`, `offset/`, `slope/`), and `ImageBaseDir` is the default directory where CISSCAL will look for images to be calibrated. Note that because CISSCAL saves and re-uses simulated dark current files, the user will need to have write permission to the `CalibrationBaseDir` directory.

It is also a good idea to add the CISSCAL directory to your `IDL_PATH` system variable if you wish to call CISSCAL from a directory other than that specified in `CisscalDir`. If the `cs`h or `tc`sh shells are being used, this may be accomplished by adding the following line to your shell script:

```
setenv IDL_PATH {$IDL_PATH}:~/cisscal3_8
```

CISSCAL has been tested under IDL version 5.5 and newer, in Linux and Unix environments. It may not be compatible with other operating systems such as Windows.

Basic Layout

CISSCAL is launched by typing "`@cisscal`" at the IDL prompt, or "`idl cisscal`" from the terminal prompt. Doing so will launch the main graphics widget.

The CISSCAL menu bar contains several pull-down menus. The menu buttons are labeled "File," "Image," "Tools," "Options," and "Batch Mode." Below the menu bar is a text window which logs to the GUI any messages generated by CISSCAL. Image label display and output logging options can be adjusted from the "Options" menu, as discussed below. To the right of the log window is a list of calibration options to be set by the user. These options are executed in the order listed, and can be toggled on and off using the buttons in the "On/Off" column. The calibration options follow the steps

described in Section 4.2, namely:

- LUT conversion
- Bitweight correction
- Subtract bias/Remove 2-Hz noise
- Subtract dark
- A-B pixel pairs
- Linearize
- Flatfield
- Convert DN to flux
- Correction factors
- Geometric correction

With the exception of the geometric correction, the default value for each of these options is "ON."

To the immediate left of the calibration option names is another column of buttons which toggles the "Option Parameters" field for that option. These parameters appear in the lower right of the GUI. Some calibration options will not have any user-definable parameters, while others have many, and can be controlled quite specifically. See the section on each calibration option for details.

The general order of operations in CISSCAL is as follows: 1) read in an image file, 2) select the desired calibration options, 3) select the desired parameters for each calibration option, 4) go to "Calibrate Image" under the image menu to execute the desired calibration steps in the order listed, 5) save the output image to a real-format VICAR image file.

Default Options File

The default calibration options used by CISSCAL are specified in the `cisscal_default_options.txt` file included with the CISSCAL distribution. This file is user-editable, though care should be taken not to corrupt the formatting.

The default options file consists of a list of 27 keywords and their values, separated by a colon. Allowed values are given in the file itself, and should be relatively straightforward for users familiar with the basic calibration options. Only the first 24 keywords are calibration options per se. In addition there are three output options: the default output filename suffix, which will replace the input filename suffix when the output VICAR file is written (this is initially set to “.IMG.cal,” but alternatives such as “_cal.IMG” may also be used); the output saturated pixel value; and the output missing pixel value. The latter two options may be set to a specific numerical value (like 0), or undefined (by setting to “NAN”), or left blank to leave the pixel value as-is. Initially, the output saturated pixel value is set to no change, and the output missing pixel value is NaN.

For a key to the non-binary (“Y” or “N”) options file settings, see the description of the CISSCAL command line mode, below.

Pull-Down Menus

The entries in the pull-down menus across the top of the main CISSCAL widget are as follows, with the menu name given in italics:

File: Open

This menu item opens a VICAR image file and stores its image data, binary header, and label information in memory as an image object. Two ancillary arrays are also generated at this time: a missing pixel array, and, if the VICAR image is in integer format (that is, presumed to be raw and uncalibrated), a saturated pixel array as well. Once an image has been read into memory, the user can operate on it with the selected calibration options by choosing "Calibrate Image" from the Image menu.

While only one image may be read into memory at a time, CISSCAL does have a limited batch mode. In this mode the "Open" button is not used at all, and images are read in sequentially and

individually.

File: Save Image

Selecting the "Save Image" item prompts the user to name an output file to which the currently defined image will be saved in VICAR format. Saturated and missing pixel values are set at this time according to the user's preferences in the default options file. The output VICAR file format is identical to the input file format except for three main differences: 1) the output image is in real (floating-point) format instead of integer; 2) the output image lacks a binary line prefix (that is, overclocked and extended pixels and other related data); and 3) the output VICAR label includes a list of appended keywords with a record of calibration tasks performed. See Table 9 for a complete list of these keywords. The last keyword added to the VICAR label is CALIBRATION_STAGE, and is used by CISSCAL to determine the last calibration step performed, so that a user may process a previously-calibrated image without repeating steps.

Note that, beginning with CISSCAL version 3.7, the output image does include the binary telemetry header, which is copied directly from the input image.

File: Save Calibration Options File

This menu item allows the user to save the current calibration options, including On/Off status as well as individual option parameters, to a binary-format file that can easily be read in at a later time.

File: Load Calibration Options File

This menu item loads a previously-saved calibration options file.

Keyword	Value
TASK	'CISSCAL 3.7'
USER	[username]
DAT_TIM	'Sat Jun 1 20:45:00 2014'
DATA_CONVERSION_TEXT	'Converted from 8 to 12 bits'
UNEVEN_BIT_WEIGHT_CORRECTION_FLAG	1
BIAS_SUBTRACTION_TEXT	Subtracted bias (overclocked pixels); Subtracted 2 Hz noise (overclocked pixels)
DARK_CURRENT_CORRECTION_TYPE	'Interpolation method'
DARK_CURRENT_PARAM_FILE	'nac_dark_parameters.xdr'
AB_PIXEL_CORRECTION_FLAG	1
NONLINEARITY_CORRECTION_FLAG	1
FLATFIELD_CORRECTION_FLAG	1
SLOPE_FILE_NAME	'SLOPE258.IMG'
FLAT_FIELD_FILE_NAME	'nac_dustring_1559626712.img; nac_dustring_aboff.img; nac_mottle_1444733393.tif'
GAIN_CORRECTION	'30.27 e-/DN'
EXPOSURE_CORRECTION_FLAG	1
SHUTTER_OFFSET_FILE_NAME	'nacfm_so_p5.img'
EXPOSURE_OFFSET	'2.75 ms'
RADIOMETRIC_CORRECTION_TEXT	'Divided by area/solid angle; Divided by efficiency factor of 86.3521; Multiplied by absolute correction factor of 1.02139'
UNITS	'1/F'
GEOMETRY_PROJECTION_TYPE	'LINEARIZED'
CALIBRATION_STAGE	'11'

Table 9: Output VICAR keywords appended by CISSCAL 3.7 and example values for a typical fully-calibrated image.

File: Write Saturated Pixel File

The user will be prompted to specify an output filename (default: "<image_name>.sat") for writing the saturated pixel array to an output image file. This array has 1s for pixels with DN values of 255 (8-bit image) or 4096 (12-bit image), and 0s otherwise. The output file will be a real-format

VICAR image with the same label as the normal output image file.

File: Write Dark Sky Mask File

The user will be prompted to specify an output filename (default: “<image_name>.mask”) for writing the masked pixels array to an output image file. This array is generated by the 2-Hz noise removal algorithm (image mean method), which uses it as a way to exclude non-dark sky pixels from the noise estimate. Pixel values of 1 are considered masked and 0s are unmasked. The output file will be a real-format VICAR image with the same label as the normal output image file.

Note that the dark sky mask created by the 2-Hz algorithm is based on a simple threshold method, which may not yield adequate results in all cases, particularly if there are detailed surface or ring features in the image. In such cases an external mask file may be created by the user and read in from the 2-Hz noise parameter window.

File: Write Missing Pixel File

The user will be prompted to specify an output filename (default: “<input image name>.miss”) for writing the missing pixel array to an output image file. "Missing pixel" is here defined as one where $DN = 0$ in a contiguous block, and the array has 1s for missing pixels, and 0s otherwise. This array is generated whenever an image file is first read into CISSCAL, and is used by the 2-Hz noise removal algorithm as a way to ignore these pixels when constructing the 2-Hz signal to be removed. The missing pixel file will be a real-format VICAR image with the same label and dimensions as the output image file.

File: Read Missing Pixel File

If the user is processing a previously-calibrated image file, it may be necessary to read in a

missing pixel file created during the previous CISSCAL session. The array read in using this menu option will replace the missing pixel array that was automatically generated upon first reading in the image file.

File: Quit

Self-explanatory.

Image: Calibrate Image

Selecting this all-important menu item causes CISSCAL to execute all the currently-selected calibration options under the calibration option listing, with their associated option parameters. Note that in batch mode (see below), this button will not be used at all.

Image: View Image and View Sliding Image

These are simple image display utilities that allow one to display the image array with a variety of different scalings: linear, logarithmic, square root, histogram-equalized, and range-stretched.

The "View Sliding Image" option displays the image array in a smaller, two-sided graphics widget, with the complete image on the left, shrunk to fit the window, and a full-resolution version of the image in the right window with vertical and horizontal scrollbars.

Tools: Histogram and Histogram (Log)

Displays a histogram of pixel values. This plot can be output to a postscript file if desired (file name = "<input file name>_hist.ps").

Tools: Profiles and Average Profiles

Plots horizontal and vertical profiles (ie. "slices") of image pixel values across a given line or down a sample column. The line and sample can be specified with the cursor, otherwise it defaults to the center pixel line and sample. "Average Profiles" simply plots the average of all horizontal and vertical profiles. This plot can be output to a postscript file if desired (file name = "<input file name>_pro.ps" or "_ave.ps").

Tools: Inspect Pixel Values

Displays image and allows user to view pixel values and line/sample by moving the cursor over regions of interest.

Tools: Plot Overclocks

Plots the overclocked pixels as recorded in the binary line prefix as a function of image line number. These "pixels" do not contain image data, but are essentially a measure of the bias level during image readout, and are used by CISSCAL in the bias subtraction and 2 Hz noise removal step. (Note that the extended pixel array, also in the binary line prefix, is corrupted in the NAC, and therefore not used for calibration.)

Options: Show VICAR label

Displays the contents of the input VICAR image label in a table to the right of the main CISSCAL window.

Options: Log to...

The user may choose to log all messages to one of four locations: the GUI log window (GUI),

the IDL terminal window (stdout), the IDL error system variable (stderr), or a user-specified log file (file). If the latter, the specified file is opened in append mode, written to, and closed every time a new message is written, so that it is continually updated during the CISSCAL session.

Options: Set Log Level...

There are three "log levels" from which the user may choose: 0 (no messages logged), 1 (standard message logging - the default setting), 2 (all messages logged).

Batch Mode

In batch mode, the user selects the calibration options and option parameters he or she desires, as usual, but then clicks on the "Batch Mode" menu item instead of "Calibrate Image." Doing so will bring up a separate dialog from which the user sets the following batch options:

- Input Directory: directory containing images to be calibrated
- Input Filter/File List: see description below
- Output Directory: where calibrated images will be written
- Output Filename Extension: file suffix to replace that of the input image filename
- Dark Subtraction Options: allows user to specify a list of dark files corresponding to each image (only used in lieu of the interpolation dark model)

If the "Input Filter" toggle is selected, the user can enter a regular expression to designate which files in the directory are to be calibrated. The default is *.IMG, or all image files. Another example would be N* or W*, i.e. all NACs or WACs.

Alternatively, the user can specify use of a file list. This is just an ASCII text file containing a single-columned list of the names of images to be processed. If this file is located in the Input Directory, then full image path names are not necessary. A batch file list can be easily constructed using a command such as:

```
> ls -l *.img > batch.txt
```

There are several limitations that come with working in batch mode. For one, there is currently no way to save ancillary image data such as the missing/masked/saturated pixels array. But more generally, some images may just require more individual attention than others, and a particular set of calibration options suitable for one image may not be suitable for all.

Calibration Options

Here we describe the individual calibration options and their user-definable parameters. Not all of the calibration steps apply to both cameras and all camera modes; in some cases this is by design, and in some cases it's simply due to incompleteness in the calibration analysis. In situations where a given option or parameter does not apply to the currently defined image, the step will simply be skipped (and a relevant message logged). CISSCAL is more-or-less self-contained in the sense that the image itself tells it whether a certain calibration step should be executed. This also means that, when in doubt, it usually doesn't hurt to leave a given option “ON,” and let the software decide whether or not it needs to be performed.

It may be helpful in getting the most out of CISSCAL to think in terms of its object-oriented data structure. Whenever an image is read into CISSCAL, it is stored as an array within an Image Object structure, of which only one can be defined at a time. This structure contains the image itself along with ancillary data such as the overlocked pixel array, binary line prefix, image label, missing and saturated pixel arrays, and the method functions that operate on the Image Object.

Each calibration option may be represented by single or multiple method functions. The “Convert DN to flux” option, for instance, is comprised of four separate method functions, whereas “LUT conversion” is a single function. Additional method functions include the routines for reading and writing image arrays and displaying or plotting image values.

The calibration options and their various parameters are stored in a data structure that is entirely separate from the image object. Thus, when you read in a new image it will not have an effect on the calibration options, and vice-versa.

The steps discussed below are always performed in the same order that they are listed here, and in the main CISSCAL GUI window, according to the calibration order of operations described in Section 4.2.

LUT conversion

This option applies only to images with a `DATA_CONVERSION_TYPE` of “TABLE.” For these images, a reverse-lookup table will be applied to convert the pixel values from an 8-bit range (0-255 DNs) to a 12-bit range (0-4096 DNs). The lookup table in question, shown in Figure 3, is hard-coded into CISSCAL, and is also available in the `lut.tab` file included in the `lut/` subdirectory of the calibration support directory.

This algorithm has no user-definable parameters.

Bit-weight correction

This algorithm corrects for uneven bit-weighting, as described in the ground calibration report and Section 4.3 above. It is not applied to LOSSY-compressed or TABLE-encoded images, due to loss of information in the compression/encoding processes. The bitweight corrections are a function of camera, gain state and temperature, and the files used to correct for it can be found in `CalibrationBaseDir/bitweight/`.

This algorithm has no user-definable parameters.

Subtract bias/2-Hz noise

The “bias” is the zero-exposure DN level of the CCD chip, and subtracting the bias is a standard CCD image calibration procedure. Cassini ISS images also suffer from a bothersome type of coherent noise that results in a horizontal banding pattern across the image. This noise is introduced during image readout, and has been found to have two peaks in its power spectrum near 2 Hz. CISSCAL treats this “2 Hz noise” as part of the bias level, and so handles both in the same calibration step. No model has yet been derived that reproduces the 2-Hz signal with exact accuracy, so removal of this noise must be accomplished by using the image itself to determine the noise component.

There are three possible options for handling the bias and 2 Hz noise. The “Constant Bias” option simply ignores any 2 Hz noise and subtracts a constant value given by the `BIAS_STRIP_MEAN` found in the image header. (The `BIAS_STRIP_MEAN` is, as the name suggests, simply the average of the overclocked pixels.) The “Overclocked Pixels” option makes use of the full overclocked pixel array in the binary line prefix to subtract a line-dependent bias level that matches the 2 Hz fluctuation in the image array. The third method, “Image Mean,” uses a combination of the `BIAS_STRIP_MEAN` and dark sky pixels in the image itself to model the 2-Hz pattern.

Determining the best method to use for a particular image is not automatic, as it is a function of both the camera settings (specifically the summation, gain, and data conversion modes) and the image scene. As discussed in Section 4.3, the 2 Hz noise in 2x2 and 4x4 summed images manifests as a complicated diagonal pattern rather than a coherent horizontal banding, such that the Overclocked Pixels and Image Mean methods are not appropriate; a constant bias should instead be used. Additionally, in TABLE-encoded images, the 2-Hz noise is at a low enough level that the quantization of the TABLE encoding makes an accurate reconstruction of the 2-Hz contribution quite difficult; a constant bias subtraction is usually the best option for these images as well.

For un-summed 12-bit or 8LSB images, either the Overclocked Pixels or Image Mean method

should be used, with the preferred option depending on the image scene. Specifically, for images containing one or more continuous patches of dark sky spanning the entire vertical extent of the field of view (that is, dark sky is present in a significant portion of all image lines), the Image Mean is usually (but not always) best; if no such vertical patches of dark sky are present, use of the Overclocked Pixels will be necessary.

In the Overclocked Pixels method, a linear fit to the overclocked pixel array is first subtracted out to account for the constant bias offset. Next, the overclocked pixels are smoothed, then filtered to remove high-frequency random noise components, then filtered again to remove slow-varying and DC-offset components. The resulting one-dimensional line-dependent array is then subtracted from each sample column.

The Image Mean algorithm operates on the same principle as the Overclocked Pixels method, but requires a bit more attention on the part of the user. In order for the software to be able to ignore all bright areas of the image, a dark sky mask must be used. The mask format is binary, with 0 indicating the dark background or sky pixels that will be used in the calculation, and 1 indicating masked pixels. This mask can either be created externally by the user, saved as a VICAR image, and read into CISSCAL using the “Choose mask file” option, or it can be created automatically by using the “Auto mask” feature. If “Auto mask” is selected, CISSCAL will create a mask based on the threshold and pixel range variables supplied by the user. Threshold is simply the DN value below which the image is assumed to be dark sky - if no threshold is supplied (i.e. the parameter field is left as 0.0) then a threshold will be automatically calculated. Note, however, that this automatic threshold calculation assumes the image is a star field, so for images with large non-stellar bright regions, the calculated threshold may not reflect the ideal cut-off. The pixel range parameter specifies how much smoothing is performed on the image before creating the mask. A default value of 9.0 is good for most cases. Higher values will create a more conservative mask (that is, more of the image will be masked, and less will be

considered dark sky). Once a mask file has been read in or the “Auto Mask” parameters chosen, the user can display the mask by clicking the “Show Mask” button. In most cases, judicious use of the image display and “Show Mask” features should allow the user to establish the appropriate values for threshold and pixel range.

When calibration is performed, the 2-Hz removal algorithm constructs an approximation of the 2-Hz signal as a function of line by replacing the masked and missing pixel regions by an average of the unmasked pixels in the same line, applying a smoothing filter, and then taking a median over the sample direction. This signal is then filtered the same way as in the overlocked pixel method, with a low-pass filter to remove high-frequency random noise, and a high-pass filter to remove the DC-offset/bias component, and the result is subtracted out.

Readout discontinuities in slow-readout images can cause a sudden jump in the dark/RBI level, requiring the filtering algorithm to treat the pre- and post-discontinuity regions separately. CISSCAL automatically identifies and handles these types of images appropriately.

Subtract dark

Dark current in the Cassini ISS cameras is comprised both of the traditional, slowly-increasing kind that we typically associate with dark current, as well as an effect called residual bulk image (RBI). This effect is also time-dependent, but results from the leaking of electrons into the pixel wells from whatever image was on the CCD previous to the current exposure. It is because of this RBI effect that the ISS cameras typically engage a “light-flood” step before every exposure - this ensures that the RBI pattern will at least be uniform across the array.

There are currently two possible methods for removing the dark current. The most straightforward is to simply read in an external VICAR-format dark file created by the user from actual in-flight dark images, or with whatever method he or she feels is appropriate.

However, because the dark current is a function of how long a given pixel sits on the CCD chip before being read out, and because this in turn is a function of a number of camera parameters including compression mode, compression ratio, summation, read-out index, telemetry rate, etc., it is simply impossible to take enough in-flight dark images to represent the full range of possibilities. The alternative, therefore, is to create a dark image by using a model.

As described in Section 4.3, the current best method for removing dark current is the interpolation method, which is now implemented for all camera modes. The interpolation method requires first defining a parameter file containing the average RBI/dark DN count for each pixel as a function of exposure time. Eight discreet times are used, covering the entire exposure range: 0, 10, 32, 100, 220, 320, 460, and 1200 seconds. The parameter file is not created within CISSCAL, but by an external code written by Bob West at JPL (robert.a.west@jpl.nasa.gov). NAC and WAC parameter files are included with the software, and can be found in the darkcurrent/ subdirectory of the calibration support directory.

Analysis of ISS dark images over the course of the mission indicates that changes in the dark pattern over time are mainly limited to the effects of what we refer to as “hot pixels”: pixels whose DN values are erroneously high compared to that expected from the dark current/RBI rate. These pixels are a concern for image calibration, as the very brightest ones can easily saturate with large exposure times, and can produce noticeable vertical streaks in the dark/RBI pattern at readout. The overall number of hot pixels is increasing over time, with some remaining as permanent features of the CCD, and others turning off and on seemingly randomly. For this reason, the current best way to handle them is to use dark calibration images acquired over the course of the mission to construct individual “snapshots” of the pixels for as many different times as we have adequate data. CISSCAL is then able to update each simulated dark with the hot pixel snapshot closest in time to the image being calibrated.

Beginning with CISSCAL 3.7, the darkcurrent/ directory contains several of these “snapshots”

in the form of text files, each containing a list of hot pixels (identified by sample, line) and corresponding electron counts at a range of exposure times. The time of the source observation is denoted by the subscript in the filename (in YYDOY format). All darks created using a given hot pixel file will be placed in a subdirectory within darkcurrent/ called “darks<YYDOY>/”. For this reason, it is important that the user has proper write permissions to these directories. CISSCAL will always look here for a dark file that matches the parameters of the current image. If it does not find one, it will construct a new dark and write it to a VICAR file in the same directory.

CISSCAL 3.7 also features a new label format for the simulated dark files, and new dark parameter files that have had hot pixels, saturated pixels, and low-frequency artifacts removed.

A-B pixel pairs

This option causes CISSCAL to search for and remove the bright/dark pixel pair artifacts created by the anti-blooming mode. These pairs are only produced when anti-blooming mode is ON, so images with ANTIBLOOMING_STATE_FLAG = OFF will not need to have this step performed.

The creation of anti-blooming pixel pairs is discussed in the ground calibration report and Section 4.3. Recent analysis has found that these pairs tend to change intensity and even position with time, so a static pixel pair map (like the one included in the antibloom/ directory, which was created during ground calibration) is insufficient to identify them.

The current algorithm identifies pairs automatically. There is only one user-definable parameter: a threshold level that defines the minimum DN difference between a pixel and its neighbors that will cause it to be identified as a candidate for a pair. The default value for this threshold is 30.0, which should be adequate for most cases. Lowering the threshold can lead to more pixel pairs being identified, and vice-versa, but the effect is much more significant in some images than others.

Once pixel pairs are identified, they are replaced by an average of their horizontal neighbors.

Linearize

The linearize function corrects for non-linearity of the CCD response, as discussed in Section 4.3. The linearity correction is a function of camera and gain state.

This algorithm has no user-definable parameters.

Flatfield

This step is actually comprised of two separate steps: dust ring removal and flatfield removal. Technically, there is no difference - any dust ring signatures in the image are part of the flatfield - but since we are using flatfields derived from ground calibration data (called “slope files”), and since components of the instrument have likely shifted in flight since then, we need to make additional corrections if we wish to accurately represent the current flatfield of a given filter combination.

As of CISSCAL version 3.8, the dust ring removal code applies several corrections to the NAC, and none to the WAC (which, for the most part, does not require any). The NAC corrections are for 1) rings measured during Venus flyby and using Titan atmosphere images taken over the course of the mission (these features are cumulative over time, and have been documented up through rev 225); 2) a correction to the preexisting bright ring located at sample=887, line=388, which only applies to images where anti-blooming mode is set to OFF; and 3) the dividing out of a “mottle map.” The mottle map is essentially a residual flatfield (exhibiting a pattern that could be described as “mottled”), and may be an artifact introduced by the “haze anomaly” event in C27 (see Section 3.6), or during the subsequent decontamination process. The strength of the mottle pattern appears to be somewhat dependent on filter wavelength.

The flatfield removal algorithm simply reads in the slope file appropriate for the given filter combination and temperature, normalizes it to the average of its inner 400x400 pixels, and then divides

it out of the input image. Dust ring and mottle corrections are then applied where appropriate.

The slope files are contained in the slope/ subdirectory of the calibration support directory, and their filter and temperature parameters are identified by the slope_db.tab database files.

The flatfield algorithm has no user-definable parameters.

Convert DN to flux

This calibration option is actually comprised of several separate calibration routines; specifically those encapsulating steps 9) and 10) in Section 4.2. In order of execution, these are:

- DNtoElectrons - multiplies image by appropriate gain to yield electrons.
- DivideByExpoT - divides by exposure time, correcting for shutter offset effect. This effect is due to the opening and closing of the shutter mechanism. The shutter offset correction is a function of camera and temperature, and the relevant data files can be found in the offset/ subdirectory of the calibration support directory.
- DivideByAreaPixel - divides image by optics area and solid angle. (NAC has area = 264.84 cm², solid angle = 3.6 x 10⁻¹¹; WAC has area = 29.32 cm², solid angle = 3.6 x 10⁻⁹.) In I/F mode (discussed below), solid angle is not divided out, as the input flux spectrum - typically solar - is assumed to be integrated over the source.
- DivideByEfficiency - divides image by the system transmission $T_0(\lambda)T_1(\lambda)T_2(\lambda)Q(\lambda)$ integrated over wavelength for the given filter combination. In the current implementation of the code, the quantum efficiency determined from ground calibration is also multiplied by a QE correction spectrum, $Q_{\text{corr}}(\lambda)$ (contained in the files nac_qe_correction.tab and wac_qe_correction.tab in the correction/ subdirectory), derived from absolute calibration.

The main decision the user must make in converting to flux units is whether to simply convert to standard (cgs) intensity units (phot/cm²/s/nm/ster), or to normalize the intensity in terms of some other flux. Typically this other flux will be the solar flux at the distance of the target from the sun, and the result is called I/F.

In I/F mode, the user will need to decide whether to normalize the image intensity in terms of the solar flux at a given solar radius, or in terms of some other user-defined flux. If the former, which is the usual approach, you must specify the solar distance by clicking on either “Jupiter” or “Saturn,” (which causes the distance to be calculated automatically for the image time given in the image label) or by inputting your own.

If you choose instead to normalize the image intensity in terms of a user-defined flux - essentially replacing the F in I/F with something other than F_{solar} - select “User-input” under I/F options, and click on the “Browse” button to be prompted for a file. The user-input flux file should be ASCII text and contain wavelength-flux pairs where wavelength is in nanometers over the range of the instrument (approximately 200 - 1100 nm), flux is in photons/cm²/s/nm (/ster if not integrated over object), and the wavelength and flux values are separated either by spaces or tabs. The top of the file must have a “\begindata” line, above which there can be a header of any length. For an example of this format (with slightly different units), the user should look at the “solarflux.dat” file supplied with the calibration support directory.

As of version 3.7 of the code, the user may toggle each flux conversion step individually. This may be useful if one wishes to calibrate in units of electrons, for example.

Correction factors

Because there is always going to be a slight discrepancy between the flux we expect given the integrated system transmission function, and what is actually observed, filter-specific correction factors have been derived to force the observation and theory to match. This calibration option simply divides the image array by these correction values, which are contained in a database file called `correctionfactors_qecorr.tab`, located in the `correction/` directory.

This algorithm has no user-definable parameters.

The correction factors were calculated starting with component calibrations for the transmission of the optics, filters and quantum efficiency, shutter performance and gain. Of these, uncertainty in the quantum efficiency (approximately 20%) dominates the uncertainty budget. In flight we rely primarily on photometric standard stars with ancillary information from calibrated observations of icy satellites, Saturn, and Saturn's rings.

To calculate the NAC correction factors used by CISSCAL 3.6, we made use of three different data sets: Vega images taken specifically for ISS calibration between S17 and S41, low-phase Enceladus images (confined to a narrow range in subspacecraft latitude and longitude to minimize systematic brightness differences) and a set of UV-bright stars (the “Bright 19”) taken during cruise. All images were processed using CISSCAL, cosmic rays removed, and photometry performed in IDL. For the Enceladus images, both a phase correction and a wide-field PSF correction were applied. Filter-specific correction factors were then calculated by taking the ratio of the observed integrated stellar fluxes to the expected values.

For the WAC, we used images of Vega, 77 and 78 Tau, and HR996 from S17 through S41, and tied these star data to the NAC results using BOTSIM images taken using identical filter combinations where available. After performing calibration, cosmic ray removal and photometry as before, the resulting correction factors for all targets were normalized to Vega in the IR, where overlap was seen to be most consistent. We then used a weighted average to calculate final factors. As a final step we applied a linear correction to both the NAC and WAC QE correction functions as a way to “split the difference” between the absolute normalization level of each, while keeping the overall NAC/WAC ratios the same.

The uncertainty of stellar fluxes is on the order of 10%, so this is the uncertainty we expect to achieve when calibration is complete. Much more detail about the absolute calibration is provided in

the end-to-end example in Section 5.

As of CISSCAL 3.7, a separate absolute calibration option (discussed in Section 4.3) has been added for Jupiter. Note that this option is separate from the Jupiter option when specifying the solar distance in the DN-to-flux conversion step.

Geometric correction

This option performs a 2-D distortion transformation on the image array, to account for the geometric distortion introduced by the optical system. Fortunately, distortion in the ISS cameras is quite insignificant in most cases: the NAC has almost no distortion at all, and while the WAC has some noticeable pin-cushion distortion, it yields an image translation of only about a pixel at the image corners. Still, correcting for this distortion may be desirable in cases where calculating accurate positions of features is crucial.

This algorithm has no user-definable parameters, and is turned off by default.

Running CISSCAL from the IDL Command Line

This software distribution contains a supplemental IDL routine, `cisscal_cl.pro`, which allows the user to run CISSCAL from the command line. As with the GUI version of CISSCAL, the input and output files are in VICAR format (as opposed to, say, IDL image arrays). This program can be used to calibrate individual images or lists of images, the latter of which may be specified either as a regular expression, an array of filename strings, or alternatively, by using the “readlist” keyword, as a single-column text file containing a list of filenames. In addition the user may designate a calibration options file to be read, allowing for the use of multiple options files suitable for different targets or image types, for example. Certain calibration options may also be specified explicitly by keyword, overriding the calibration options file settings.

The calling sequence, with optional keywords in brackets, is:

```
cisscal_cl,file,[optionfile=?, readlist=?,batchdir=?, outputdir=?,  
suffix=?, default=?, bias=?, twohz=?, im_threshold=?, im_pixrange=?,  
flux=?, geom=?, spec=?, mask=?]
```

where:

- `file` = string containing name of single image, or regular expression ('N*' for example) or comma-delineated list of images to be calibrated.
- `optionfile` = name of options file to be used; if left blank, will be automatically set to `CisscalDir/cisscal_default_options.txt`, or user can override `optionfile` with `/default` keyword.

The following key describes the possible settings for non-binary (“Y” or “N”) options:

Bias subtraction options:

“BSM” = bias strip mean (default)

“OC” = overclocked pixels

2 hz removal options:

“IM” = image mean

“OC” = overclocked pixels (default)

Flux conversion options:

“I” = intensity units

“IOF” = I/F (default)

distance from sun = J (Jupiter), S (Saturn; default), AU (manual input in astronomical units)

- `readlist` = string containing name of file with single-column list of images to be calibrated; overrides `file` option.
- `batchdir` = input directory, if not ''
- `outputdir` = output directory, if not ''
- `suffix` = extension to add to calibrated image filenames; default is '.IMG.cal'
- `/default` - uses default parameters

In addition, the following keyword options to `cisscal_cl.pro` can be used to override the settings in the

optionfile:

- bias = BSM or OC or OFF
- twohz = IM or OC or OFF
 - im_threshold = (auto = 0.0)
 - im_pixrange = (default = 9.0)
- flux = I or IOF or OFF
- /geom - (turns on geometric correction)
- spec = text variable set to spectrum file for IOF mode
- mask = text variable set to mask file for 2hz noise removal

For example, the command:

```
IDL> cisscal_cl, '*.IMG', bias='BSM', twohz='IM'
```

will calibrate all images (files ending with “.IMG”) in the current directory using the default calibration options for all calibration steps except: 1) BIAS_STRIP_MEAN used for bias subtraction (bias='BSM'); and 2) Image Mean used for 2 Hz removal (twohz='IM').

Make Polar Image

Also included with CISSCAL (version 3.7 and up) is the standalone code `make_polar_image.pro`, which can be used to generate useful quantities from polarized images. This function ingests registered or unregistered, calibrated, polarized ISS images. Input images must be in VICAR format, and specified in sets of two or three, with filters consisting of either 1) P0, P60 and P120, 2) IRP0 and IRP90, or 3) CLR and IRP0, with matching colors in the opposite filter position. The order of input images is irrelevant, only that they constitute a matching set. The code includes the transmissions parallel and perpendicular to the polarizing axis of the polarizer, derived from observations. These are functions of wavelength, and different passbands have different effective

parallel and perpendicular transmissions. As of CISSCAL version 3.7, the following filter combinations are supported:

- WAC IRP0/CLR or IRP0/IRP90 + MT2, CB2, MT3, CB3 (NAC currently uses WAC values for these filters)
- NAC P0/P60/P120 + UV3, BL2, GRN, MT1, CB1, MT2, CB2

The returned value is a structure defined as follows: if the input includes two file names (consisting of filters paired with either [IRP0, IRP90] or [IRP0, CLR]) the function returns intensity I (same units as the original, and indicated by the tag name "intensity") and the Stokes parameter Q defined with respect to IRP0 (tag name: "Q"). If input consists of three files (containing filters paired with P0, P60, P120) the returned structure contains intensity, polarization (from 0.0 to 1.0; tag name: "polarization"), and theta, the angle of the electric vector relative to the camera Y axis (+/-90 degrees; tag name: "theta").

The code is invoked as follows:

```
result = make_polar_image(file1,file2,[file3],[/align],[/flip],$
                        [outfile=outfile],[silent=silent])
```

Set the /align keyword to co-register the images via cross-correlation if they have not already been co-registered. Set the /flip keyword to flip the output arrays about the x (sample) axis. Set the /silent keyword to suppress warning and info messages (but not error messages). Set the outfile keyword equal to the filename of a VICAR-format file containing the calculated intensity. This file will be given the same label as the first input file to track calibration history.

For more information on the polarization code implementation, see Section 4.3, above.

5. End-to-End Example: Absolute Flux Calibration (West et. al., 2010)

5.1 Introduction and Setup

In this section we provide a detailed example describing how ISS data can be acquired, calibrated, and used for scientific analysis. The analysis in question is the derivation of the absolute calibration correction, as summarized in West et. al. (2010), which establishes the absolute flux and relative color corrections for each filter combination used by ISS. This analysis uses both star and satellite targets, and a fairly large number of images of each, so automation of computationally-intensive tasks is a must.

Users who have read up to this point in the Data User's Guide should have all of the information they need to understand the basis of the absolute calibration, and to begin downloading and processing ISS data. At minimum, users should read Sections 4.2, 4.3 and 4.4 before they proceed. The “Absolute Calibration” subsection of Section 4.3 in particular contains necessary context and a basic outline for the procedure that we'll describe here.

As explained in Section 4.4 and elsewhere, the CISSCAL software package is written in IDL, and so access to and familiarity with that language are assumed in the pages that follow. Basic file navigation, text editing, and use of a terminal command line in Linux/Unix will also be assumed, though CISSCAL has been successfully tested on Macintosh systems as well. (Compatibility with Windows has not been determined.)

Before beginning, the user will also need to download the appropriate calibration software and data files from the COISS_0011 Calibration volume, available from PDS here:

<http://pds-imaging.jpl.nasa.gov/volumes/iss.html>

Note that the analysis described in this end-to-end example provides the basis for the absolute

flux correction in the latest version of CISSCAL (version 3.8). This means that in order to properly reproduce the results of the analysis, the user must use the previous version of the software, 3.4, which is available on the initial release of the COISS_0011 volume. Specifically, the data files for CISSCAL 3.4 can be accessed directly at:

http://pds-imaging.jpl.nasa.gov/data/cassini/cassini_orbiter/coiss_0011/calib/calib.tar.gz

and the CISSCAL software files are likewise available in gzipped tar archive form here:

http://pds-imaging.jpl.nasa.gov/data/cassini/cassini_orbiter/coiss_0011/extras/cisscal.tar.gz

The latest versions of these files, recommended for general purpose ISS calibration, can be found at the following URLs:

http://pds-imaging.jpl.nasa.gov/data/cassini/cassini_orbiter/coiss_0011_v3/calib/calib.tar.gz

http://pds-imaging.jpl.nasa.gov/data/cassini/cassini_orbiter/coiss_0011_v3/extras/cisscal.tar.gz

See Section 4.4 for instructions on how to install and set up CISSCAL. The calib/ directory should be copied to a local working directory to which the user has read and write permissions. Note that this directory is fairly large, at nearly 700 MB in size in its compressed form.

Figure 14 shows the results of the previous absolute flux calibration, completed in December of 2004, and on which the CISSCAL 3.4 absolute flux correction is based. (The correction used previous to CISSCAL 3.4 is indicated by the dashed line.) This analysis used Vega as the sole standard, with a comparison spectrum taken from Bohlin & Gilliland (2004). These Vega images had relatively short exposure times - as low as 5 ms, the shortest exposure setting available - and thus were particularly helpful in identifying a proper shutter offset time (~2.85 ms for the NAC), which has since been incorporated into the calibration software. However, these short exposures also contribute significant shutter uncertainty to the results.




Figure 14: Earlier Vega-derived absolute corrections for the NAC. The smooth curves correspond to the QE correction function used in CISSCAL 3.4 (solid) and previous to that (dashed). Residual offsets are recorded as filter-specific correction factors.

From the same Vega photometry we computed an absolute flux correction, which is implemented in the calibration software in two parts:

- Q_{corr} , a smooth, wavelength-dependent function which covers the entire ISS bandpass, considered to be a correction to the detector quantum efficiency function Q . (See Equation 2; data contained in `nac_qe_correction.tab` and `wac_qe_correction.tab` in the `calib/efficiency/` subdirectory.)
- $C_{f1,f2}$, filter-dependent residual correction factors. (See `correctionfactors_qecorr.tab` in the `calib/efficiency/` directory.)

The total integrated transmission of the instrument as a function of filter combination is then given as:

$$T_{tot} = C_{f1,f2} \int T_1 T_2 T_o Q Q_{corr} d\lambda \quad (14)$$

The uncertainty in the absolute calibration of the ISS cameras after the original Vega-based calibration effort was between 10-25% for the NAC and somewhat higher for the WAC, depending on filter combination. Our goal was to get these uncertainties down to around 10% for all filters, which is about the absolute flux uncertainty of Vega, our primary standard star.

For this analysis we approach the problem by initially leaving the Vega-derived quantum efficiency correction as-is, and calculating a new set of filter-dependent correction factors that will supercede the previous one. Small tweaks to the quantum efficiency function can then be made as a final step, as a way both to establish the overall normalization level, and to make sure that the NAC and WAC results agree with each other as well as possible across the entire instrument bandpass.

In the sections that follow we'll discuss the data requirements for our analysis and use these to construct an appropriate search strategy for navigating the PDS Planetary Image Atlas. We'll also explain the most efficient ways to download the data and propose methods for keeping track of the necessary image metadata. Next we'll guide the user through the process of image calibration, paying particular attention to the use of any settings beyond the defaults. Specialized topics covered here will include the use of mask files in the bias subtraction and 2 Hz removal routines, and use of input reference spectra for calculating intensities normalized to a known spectrum. Finally, we'll discuss the image analysis, including photometry on both point source and extended targets, and methods for weighting, combining and interpreting the results. In conclusion we'll tie these results to those we reported in West et. al. (2010).

5.2 Assembling the Data Set

Comparison Spectra

The absolute calibration analysis is constrained, first and foremost, by the spectra we have available to use as an absolute comparison. For the present analysis, we have the following comparison spectra for the NAC:

- Enceladus spectrum from Keith Noll (2008, private communication), 200 - 1150 nm, trailing hemisphere spectrum used for both leading and trailing hemispheres;
- Vega spectrum from Bohlin & Gilliland (2004), 115-2600 nm;
- 14 stars from the “Bright 19” data set of UV-bright stars (observed during a UVIS ride-along during cruise): HR6527, HR5191, HR1713, HR1790, HR1948, HR2004, HR3165, HR472, HR6175, HR2294, HR2491, HR2618, HR5267, and HR1903; all spectra longward of 320 nm taken from the Pulkovo Spectrophotometric catalog (Alekseeva et al., 1997, Vizier designation III/201) with the exception of HR1903, taken from the Southern Spectrophotometric Standards catalog (Hamuy et al., 1992/1994, Vizier designation II/179); UV portion of spectra derived from Kurucz 1993 model and normalized to match observed data.

And for the WAC:

- Vega spectrum from Bohlin & Gilliland (2004), 115-2600 nm;
- HR996 spectrum from Glushneva et. al. (1998) for 322.5-762.5 nm; Santos et. al. (2001) for near-IR (normalized to Glushneva spectrum); Heck et. al. (1984) from IUE for 115.3-320.1 nm;
- 77 Tau spectrum from Bruzual-Persson-Gunn-Stryker catalog (1983), 22.9-2560 nm;
- 78 Tau spectrum Burnashev (1985) for 320-817 nm; Glushneva et. al. (1998) for 322.5-762.5 nm; Kharitanov et. al. (1988) for 322.5-757.5 nm; Glushneva et. al. (1998) for 597.5-1082.5 nm (normalized to average of previous three); Jamar et. al. (1976) for 136-254 nm; Heck et. al.

(1984) from IUE for 115.3-320.1 nm.

All comparison spectra except for that of Vega were interpolated to a 1-angstrom grid, normalized as indicated above, averaged to combine in the overlapping regions, and the final results smoothed with a 50-angstrom window (150-angstrom for Enceladus) using the IDL 'smooth' function.

Summary of Data

The image selection criteria for each target are given below. As a general rule, for absolute calibration, we want to obtain the highest-fidelity, highest-resolution images possible. Table 9 contains a summary of keywords and search values to be used for querying the Planetary Image Atlas.

Enceladus

Enceladus images were chosen for this analysis based on the following criteria:

- All images un-summed, losslessly-compressed, 12-bit.
- Sub-spacecraft phase angle < 30 degrees.
- Target distance such that entire satellite fits within the field of view (generally > 100,000 km), and subtending at least 100 pixels in diameter (corresponding roughly to a pixel scale < 5 km/pixel).
- No corrupted data or missing lines near target.
- Excluded all data outside a narrow range in geometry (sub-spacecraft latitude < 10 deg, sub-spacecraft longitude < 250 deg) to avoid light curve variations.

All Enceladus images used in this analysis were taken with anti-blooming camera mode turned ON. As this imaging mode has been seen to cause excess noise in certain long-exposure images, we sought to minimize this problem by additionally excluding all images with fewer than 50,000 pixels on

target and with exposure times greater than 1 sec whose background noise, as measured in a horizontal strip across the image, exceeded an empirical threshold (standard deviation > 0.1).

The above selection criteria yields two sets of data, one for the leading hemisphere (sub-spacecraft longitude < 180 deg) and one for the trailing. These two data sets were then treated as a single set for the final calculation combining the results from the various targets.

Stars

All of the primary ISS star calibration observations – those with OBSERVATION_ID containing “CALSTAR” (see Appendix B) - are taken in 1x1 summation mode, 12-bit conversion mode, and with lossless compression. All of the “Bright 19” images used in this analysis were similarly un-summed and losslessly-compressed, though about half of them (HR1903, HR6527, HR5191, HR1713, HR1790, HR1948, HR2004) were imaged only in 8-bit LUT mode, adding between 1-2 percent of uncertainty to these data. Also excluded were any exposures exhibiting DN saturation of the target, as well as those with exposure times less than 40 ms, in order to minimize the effects of shutter uncertainty.

For the WAC absolute calibration analysis, all star images were un-summed, 12-bit, and losslessly-compressed. We also set the minimum exposure time somewhat lower, to 25 ms, so as to retain enough data for good image statistics while eliminating most of the contribution from shutter uncertainty.

Target	PDS Keyword	Operator	Value [units]
<i>Enceladus</i> (NAC only)	TARGET_NAME	=	ENCELADUS
	INSTRUMENT_ID	=	ISSNA
	INSTRUMENT_MODE_ID	=	FULL
	DATA_CONVERSION_TYPE	=	12BIT
	INST_CMPRS_TYPE	=	LOSSLESS
	PHASE_ANGLE	<	30 [deg]
	TARGET_DISTANCE	>	100000 [km]
	PIXEL_SCALE	<	5 [km/pixel]
	SUB_SPACECRAFT_LATITUDE	Between	-10, 10 [deg]
	SUB_SPACECRAFT_LONGITUDE	Between	80, 100 (leading); 200, 250 (trailing) [deg]
<i>Vega</i> (NAC and WAC)	OBSERVATION_ID	Contains	CALSTAR1
	INSTRUMENT_MODE_ID	=	FULL
	DATA_CONVERSION_TYPE	=	12BIT
	INST_CMPRS_TYPE	=	LOSSLESS
	EXPOSURE_DURATION	≥	40 (NAC); 25 (WAC) [msec]
<i>77/78 Tau</i> (WAC only)	OBSERVATION_ID	Contains	CALSTAR2
	INSTRUMENT_ID	=	ISSWA
	INSTRUMENT_MODE_ID	=	FULL
	DATA_CONVERSION_TYPE	=	12BIT
	INST_CMPRS_TYPE	=	LOSSLESS
	EXPOSURE_DURATION	≥	25 [msec]
<i>HR996</i> (WAC only)	OBSERVATION_ID	Contains	CALSTAR3
	INSTRUMENT_ID	=	ISSWA
	INSTRUMENT_MODE_ID	=	FULL
	DATA_CONVERSION_TYPE	=	12BIT
	INST_CMPRS_TYPE	=	LOSSLESS
	EXPOSURE_DURATION	≥	25 [msec]
<i>“Bright 19”</i> (NAC only)	OBSERVATION_ID	Contains	BRIGHT19
	INSTRUMENT_ID	=	ISSNA
	INSTRUMENT_MODE_ID	=	FULL
	INST_CMPRS_TYPE	=	LOSSLESS
	EXPOSURE_DURATION	≥	40 [msec]

Table 9: List of keywords and search values for querying the image database for data to be used in the absolute calibration analysis.

BOTSIM Data

In addition to the data above, this analysis uses a set of images taken in BOTSIM (“both simultaneous”) mode. Identification of these images using a limited image search interface like the one provided by the Planetary Image Atlas would be a challenge, due to the differing layouts of the filter wheels in each camera, which results in equivalent filter combinations in the NAC and WAC not necessarily having identical filter names; that is, the RED/CL2 combination in the NAC is equivalent to CL1/RED in the WAC. A better strategy would be to read the index.tab file (or better yet, a “master_index.txt” file constructed by concatenating all of the index.txt files from each data volume; such a cumulative index is not included on ISS archive volumes due to space constraints) into a spreadsheet program like Excel and use the sorting functionality of that software to help identify image pairs matching the search criteria. The criteria are:

- SHUTTER_MODE_ID = BOTSIM
- Equivalent filter combination in both cameras
- Saturn, satellite or ring target that fills most or all of NAC image
- WAC should be 1x1; lossless, 12-bit preferred for both cameras
- PHASE_ANGLE < 90 degrees to minimize stray light

The images we identified using these criteria are listed by IMAGE_NUMBER in Table 10, below.

Downloading Images from the Planetary Image Atlas

PDS provides an interface to all of its archived image data in the form of the Planetary Image Atlas, located at the following URL:

<http://pds-imaging.jpl.nasa.gov/search/search.html>

Section 3.10 of this document provides strategies for using the Planetary Image Atlas to access ISS

image data, and Tables 10 and 11 contain all of the search information necessary to begin assembling the absolute calibration data set using those instructions. Some additional manual intervention will be required to identify images tainted by missing lines, target pixels with saturated DN values (i.e. values of 4095 in 12-bit data), or those with excess noise as described in the Enceladus data summary above. To exactly replicate the data set used for West et. al. (2010), it will also be necessary to limit the results to data taken prior to SEQUENCE_ID = S51. Images of Enceladus containing other bright satellite targets should also be rejected.

Filter (+ clear)	IMAGE_NUMBER	Filter (+ clear)	IMAGE_NUMBER
BL1	1476822923 1560352181 1564709295	IR2	1476797731 1481723851 1481728529
CB2	1476823705 1481591508 1492334135 1564813961 1565204564 1565865278 1573436439	IR3	1476792697 1476823285
CB3	1564362917 1569440744 1573461751	IR4	1476823810
CLR	1477593992 1487302209 1491984145 1503162698 1556335260 1560312835 1567129584 1568127472 1568133373	MT2	1476818534 1480963080 1486388594 1488310912 1492338227 1492929505 1568875816 1573112809
GRN	1481592338 1492928328 1544899964 1546272412 1547824536 1560348946 1560352289 1564707457	MT3	1486213451 1486215948 1492334024 1499317635 1499318901 1564369305 1573337795
IR1	1476792541 1544899705 1544912663 1546272215	RED	1476812869 1544913047 1546272350 1546278640 1547836691 1564707395 1564709949

Table 10: List of images used in BOTSIM analysis, given by IMAGE_NUMBER (identical for NAC and WAC images taken in BOTSIM mode).

The attached image labels within the image files contain most of the metadata that we will need to calibrate and analyze the data. Some useful information, however, is not included in the image labels; specifically, any of the target geometry information calculated by Autonav (see Table 5 in Section 3 for a list of these keywords). Thus it is necessary to keep track of this data by some other method. We recommend using the Planetary Image Atlas to help with this by adding any relevant geometry keyword columns to the search results using the “Table Columns” selection list, and then using the “Download Report” option to produce a tab-delimited text file containing the chosen keyword values. Users may also want to include the “FILE_NAME” keyword to simplify any future operations that will require the matching of keyword values to the downloaded image files.

5.3 Image Calibration and Data Analysis

NAC Calibration and Photometry Procedure

Enceladus

We begin the calibration process by using our Enceladus reference spectrum to calculate filter-integrated comparison I/Fs to compare directly to the globally-averaged I/Fs we will measure in our satellite images. This is done by integrating $A_{\text{sat}} F_{\text{solar}} * T_1 * T_2 * T_o * Q * Q_{\text{corr}}$ over wavelength (where A_{sat} is the comparison albedo spectrum and F_{solar} is the solar flux spectrum given in solarflux.tab in the calib/efficiency/ subdirectory) and dividing the result by the integral of F_{solar} over the same bandpass. For convenience, the results of this calculation are provided in Table 11.

Filter Combination	Effective Wavelength (nm)	Comparison I/F
BL1/CL2	450.87	1.33
BL1/GRN	497.45	1.37
CL1/BL2	439.93	1.32
CL1/CB1	619.46	1.38
CL1/CB2	750.47	1.38
CL1/CB3	938.03	1.32
CL1/CL2	610.68	1.37
CL1/GRN	568.2	1.38
CL1/IR1	751.94	1.38
CL1/IR3	929.81	1.33
CL1/MT1	618.9	1.38
CL1/MT2	727.44	1.37
CL1/MT3	889.22	1.35
CL1/UV3	338.19	1.25
HAL/CL2	655.66	1.39
IR2/CL2	862	1.37
IR2/IR1	827.48	1.38
IR2/IR3	901.87	1.35
IR4/CL2	1001.64	1.26
IR4/IR3	996.31	1.27
RED/CL2	650.12	1.38
RED/GRN	601.09	1.38
RED/IR1	701.97	1.37
UV1/CL2	257.96	1.16
UV2/CL2	297.82	1.21
UV2/UV3	315.61	1.22

Table 11: Comparison I/F values derived from Keith Noll's Enceladus spectrum by integrating over the ISS filter bandpass. Effective filter wavelength in nanometers is also provided.

Image calibration is performed in CISSCAL, using default image calibration settings for all calibration steps with the exception of bias and 2-Hz noise removal, which require a slightly more careful approach. In images such as these, where there is a continuous strip of dark sky throughout the entire vertical extent of the field of view, the “Image Mean” method is the preferred method for isolating and removing the 2 Hz pattern, as it produces a better, less noisy average. This method is aided by the ability to use a mask file to mask out portions of the image other than the dark sky

background. However, in some images, the target is large enough in the field of view that light from the extended PSF creates a “halo,” adding a gradient to the image mean and thus making it less useful for 2 Hz subtraction.

For satellite calibration, we handle this as follows: for images in which the target satellite subtends *fewer* than about 300 pixels in diameter, we create a dark-sky mask file (one for each OBSERVATION_ID is usually adequate, unless the target size or position have changed drastically), and use this along with the “Image Mean” 2 Hz noise removal in CISSCAL. In these cases, the BIAS_STRIP_MEAN method should be used to approximate bias. Whereas, for images in which the target satellite subtends *greater* than about 300 pixels in diameter, no mask file is created, and the overclocked pixel arrays should be used for both 2-Hz noise and bias level removal.

The dark sky masks can be created using a straightforward smoothing/thresholding method. In IDL, an example of such a method using the “smooth” routine might look like:

```
IDL> mask = smooth(img,30) > thresh
IDL> mask[where(mask eq thresh)] = 0.0
IDL> mask[where(mask gt 0.0)] = 1.0
```

where *img* is the raw input image array, and *thresh* is a threshold value set to just above the image background level. Here we have chosen a 30-pixel smoothing window.

Calibration now proceeds in CISSCAL. Considering the large number of images to be calibrated, most users will want to take advantage of the batch processing option, either by selecting “Batch Mode” from the GUI menu bar, or by using the file input options available from the command line (using `cisscal_cl.pro`). Even so, it will be necessary to calibrate in several batches, one for each mask file/target geometry, since CISSCAL recognizes only one mask file at a time.

After calibration and before photometry, a boxcar-average cosmic ray removal algorithm (`acre.pro`, available from the University of Washington Astronomy IDL search page at <http://www.astro.washington.edu/deutsch/idl/>) is applied to all images, and then photometry performed

in a straightforward manner: by summing the total I/F and then dividing by the total number of pixels on the target,

$$n_{pix} = \frac{\pi r_{sat}^2}{pixscale^2} \quad (15)$$

where *pixscale* is the image PIXEL_SCALE in km/pixel. (Ideally the pixel scale would be derived from the TARGET_DISTANCE so as to correspond to the target center rather than the sub-spacecraft point, though at most this makes a fraction of a percent of difference.)

Two additional corrections were applied to the Enceladus photometry data in the West et. al. (2010) analysis. First, a phase angle correction using an Enceladus phase curve model supplied by Paul Helfenstein (unpublished), and secondly, a “lost light” correction to recover any light from the extended tail of the point spread function that has fallen outside the frame. For the latter, we created a fake image of Enceladus subtending the same number of pixels as the source image, and then convolved it with the PSF for that filter. Then we could simply calculate the fraction of the resultant flux falling outside of the camera's field of view, and add this back to the original image before summing. With the exception of only a few filters, this correction made less than a 1% difference.

Correction factors are derived from the photometry results by dividing the mean globally-averaged I/F for each filter by the comparison I/F values we calculated previously.

Vega and Bright 19 stars

As with the Enceladus data, default calibration settings should be used for all calibration steps with the exception of bias subtraction, for which we used the BIAS_STRIP_MEAN method for estimating bias, and 2-Hz noise removal, for which the “Image Mean” method with default threshold and pixel range parameters was chosen. Then, instead of calculating the resulting brightness in standard flux units, an “I/F” was computed using the known target spectrum as the “F.” The resulting star-

integrated flux could then be taken as the unitless correction factor for that image.

After calibration, a cosmic ray removal algorithm was applied to all images, and photometry performed using an increasing aperture radius up to 34 pixels centered on the star centroid, with the baseline sky value calculated as the median of the ring region from two to eight pixels outside this aperture. Aperture radii between 18 to 28 pixels were found empirically to give the most consistent results with the least amount of sensitivity to changing radius size, so the final flux value was taken to be the median of these values. A few images were found to still yield excessively noisy photometry results as a function of sampling radius, possibly indicating less-than-fully-subtracted cosmic rays or other blemishes, and thus were excluded from subsequent steps in the analysis.

Combining NAC Results

Figure 15 shows the Enceladus, Vega and Bright 19 photometry results, with all data sets individually scaled to the average of CL1/GRN for Vega (although note that the star data for that filter was not used for the final correction factor calculation due to high shutter uncertainty caused by short exposure times). Values for a given filter combination are used only if there were two or more usable images for that particular target, where here the Bright 19 stars are together considered a single “target” due the limited amount of data available.

From this plot the Vega points for CL1/BL2 and BL1/GRN were determined to be erroneously low, and those for Enceladus UV2/UV3, IR4/IR3 and IR4/CL2 erroneously high, so these were also excluded from the final calculation. Having been as diligent as possible in rejecting erroneous data, we are now able to calculate the average NAC correction factors as a function of filter combination.




Figure 15: Intermediate NAC correction factors derived from Vega, Enceladus, and the "Bright 19" stars, and normalized to the CL1/GRN value for Vega. Lines have been drawn through the filter averages for each target.

Now that we have a complete set of correction factors normalized to CL1/GRN, we can then set the absolute flux level by scaling the results such that we minimize the average correction factor offset for Vega. For this step we were slightly more inclusive in the data we used, excluding all UV points (as there was only one image per filter, and they seemed erroneously low) but including CL1/BL2, BL1/GRN, and all non-saturated, non-noisy images with exposure times > 20 ms. The resulting normalization constant of ~ 0.96 was then applied to the NAC QE correction function, with the individual filter correction factors left as-is.

WAC Calibration and Photometry Procedure

For the WAC, we performed calibration on the images in the usual way, using the default settings for CISSCAL 3.4, with the exception of using the BIAS_STRIP_MEAN for bias subtraction, and image mean with default threshold and pixel range parameters for 2-Hz noise removal. As for the NAC star images, we calibrated in terms of “I/F” where “F” in this case is the comparison star spectrum. Thus, again when we integrate over the star during photometry, we end up with a unitless flux that may be taken as the correction factor for that filter.

As before, a cosmic ray removal algorithm (acre.pro) was applied to all images after calibration. And, as before, photometry was performed using a range of aperture radii centered on the stellar centroid, and then taking the median of the results within a subset of that range. The aperture radius range of 18-28 pixels was used for Vega and HR996, but for 77/78 Tau a smaller range from 8 to 12 pixels had to be used since the two stars are so close together in the frame. This results in correction factors (i.e. normalized integrated flux values) that are slightly lower for these stars compared to the others, although since we're primarily only worried about filter-to-filter differences (with other sources setting the absolute flux normalization), and because we expect the offset to be roughly the same as a function of filter/wavelength, this shouldn't be a problem.

After normalizing all data to Vega using the overlapping filters, we find that the star-derived correction factors match up remarkably well at wavelengths greater than 800 nm, though less well for the visible and IR portion of the spectrum.

BOTSIM Analysis

In order to ensure the consistency of the WAC calibration with the NAC, we made use of BOTSIM mode images (in which the NAC and WAC are both exposed simultaneously), with both

cameras sharing the same filter combination. Only 12-bit or LUT mode images were used, and only geometries with a phase angle less than 90 degrees were considered in order to minimize stray light. While image pairs fulfilling all of these criteria are in short supply, a few exist for all filters shared by the NAC and WAC except VIO and HAL. Due to the co-alignment of the two cameras, and near exact factor of 10 difference in image scale, these BOTSIM image pairs provide a straightforward method of comparing the absolute sensitivities of the cameras by eliminating all differences in target and pointing geometry.

The BOTSIM comparison process consists of the following steps:

1. Calibrate images to I/F using standard CISSCAL settings.
2. Repair or remove from consideration any missing data, cosmic rays or other problem pixels.
3. Crop the WAC image to the central 1/10th; smooth and rebin NAC image to same size.
4. Align cropped WAC image with corresponding NAC image using a Fourier-based correlator (typical adjustment was ~1 pixel).
5. Define “good” pixels as brighter than some background threshold, where NAC and WAC images not too different, and not adjacent to any “bad” pixels.
6. Calculate NAC-to-WAC ratio for all “good” pixels, fit a line to the resulting distribution using a least absolute deviation method (`ladfit.pro` in IDL); the slope of this line can then be converted to a correction factor for one camera (in this case, the WAC) by holding constant the other camera's correction factor for the filter in question. This can then be compared against the factors derived from our other WAC calibration sources.

Note that this fitting method incorporates a y-offset which might be present in the NAC vs. WAC plots due to excess background flux or scattered light in one or both images. In general the WAC correction factors resulting from this method matched up quite well with those derived from star observations, especially at wavelengths < 700 nm.

Combining WAC Results

Rather than allowing Vega to set the absolute flux normalization level as we did for the NAC, in this case we want instead to tie the WAC flux normalization to the NAC using the correction factor values derived from the BOTSIM analysis. In other words, we want to scale our Vega correction factors to the BOTSIM correction factors as opposed to the other way around. We do this by simply calculating and dividing out an average offset among overlapping filters. The HR996 and 77/78 Tau data is then scaled to the Vega data using filters with wavelengths > 800 nm, where the correlation appears tightest.

To produce the combined average correction factors for the WAC, we weighted the data according to the relative weights in Table 12 below, with the star weights corresponding roughly to the confidence we have in their respective comparison spectra.

Filters	Vega	HR996	77 Tau	78 Tau	BOTSIM
CL1,VIO	0.3	-1	-1	-1	-1
CL1,BL1	0.3	0.05	0.1	0.05	0.5
CL1,GRN	0.3	0.05	0.1	0.05	0.5
CL1,CL2	-1	0.05	0.1	0.05	0.5
CL1,RED	0.3	0.05	0.1	0.05	0.5
CL1,HAL	0.3	0.05	0.1	0.05	-1
MT2,CL2	0.3	-1	-1	-1	0.5
CL1,IR1	-1	-1	-1	-1	0.5
CB2,CL2	0.3	0.05	0.1	0.05	0.5
IR2,IR1	0.3	0.05	0.1	0.05	-1
IR2,CL2	0.3	0.05	0.1	0.05	0.5
MT3,CL2	0.3	-1	-1	-1	0.5
IR3,CL2	0.3	0.05	0.1	0.05	0.5
CB3,CL2	0.3	-1	-1	-1	0.5
IR4,CL2	0.3	-1	-1	-1	0.5
IR5,CL2	0.3	-1	-1	-1	-1

Table 12: Weight values for combining WAC correction factors from each target. -1 corresponds to no data.

Residual QE Correction

For the final step, we apply what we will call a residual quantum efficiency correction to the QE correction function for both cameras such that the deviation of our newly-calibrated Vega data as a function of wavelength is minimized as much as possible. To do this we simply re-calibrate all of the Vega data using our newly-derived NAC and WAC correction factors and NAC quantum efficiency correction, re-do the photometry, and plot the results for both cameras together to look for any overarching trends with wavelength. The results are shown in Figure 16.




Figure 16: Residual QE correction derived from Vega using the newly-derived correction factors. The fit line is applied to the QE correction function for both cameras, ensuring that the camera-to-camera ratio remains constant with filter.

Though we expect the normalized flux values to be near unity for all filters, it is evident in the plot above that there is a general upward trend with wavelength. The resulting linear fit is then divided out of the QE correction functions for both cameras, thus preserving the BOTSIM ratios for each filter

while yielding a much-improved fit across the entire wavelength range of the instrument.

5.4 Final Results

After applying the quantum efficiency correction described above, we re-calibrate the images a final time, re-perform the aperture photometry, and re-combine the results to derive final residual correction factors. These values are contained in the `correctionfactors_qecorr.tab` file in the `calib/efficiency/` subdirectory of version 2 of the COISS_0011 volume, and, along with the `wac_qe_correction.tab` and `nac_qe_correction.tab` files, set the absolute flux correction for CISSCAL 3.6.

Figures 10 and 11 in West et. al. (2010) show the final correction factors plotted against wavelength, and Tables 4 and 5 give the standard deviations, from unity as well as from the mean value. Figures 7, 8 and 9 in that document also show corrected ISS data plotted against the reference spectra of Vega, HR2294 and Enceladus, respectively.

Appendix A: ISS Instrument Data

Filter	Numerical Name (NAC)	Numerical Name (WAC)	Science Justification
UV1	258W	--	Aerosols
UV2	298W	--	Aerosols, broad-band color
UV3	338W	--	Aerosols, broad-band color, polarization
VIO	--	420SP	Broad-band color
BL2	440W	--	Medium-band color, polarization
BL1	451W	460W	Broad-band color
GRN	568W	567W	Broad-band color
MT1	619N	--	Methane band, vertical sounding
CB1	635N	--	Two-lobed continuum for MT1
RED	650W	648W	Broad-band color
HAL	656N	656W	H-alpha/lightning
MT2	727N	--	Methane band, vertical sounding
CB2	750N	752N	Continuum for MT2
IR1	752W	742W	Broad-band color
IR2	862W	853W	Broad-band color, ring absorption band
MT3	889N	890N	Methane band, vertical sounding
CB3	938N	939N	Continuum for MT3, see through Titan haze
IR3	930W	918W	Broad-band color
IR4	1002LP	1001LP	Broad-band color
IR5	--	1028LP	Broad-band color
CL1	611W	635W	High sensitivity, combine with filter wheel 2 filters
CL2	611W	635W	High sensitivity, combine with filter wheel 1 filters
P0	617W	--	Visible polarization, 0 deg
P60	617W	--	Visible polarization, 60 deg
P120	617W	--	Visible polarization, 120 deg
IRP0	--	705W	IR polarization, 0 deg; see through Titan haze
IRP90	--	705W	IR polarization, 90 deg; see through Titan haze

Table A.1: ISS filter names and their science justification. Adapted from Porco et. al. (2004).

Camera	Filter 1	Filter 2	Central Wavelength (nm)	Bandpass FWHM (nm)	Effective Wavelength (nm)	PSF FWHM (pixels)	PSF Dynamic Range	Extended PSF?
NAC	CL1	CL2	610.675	340.056	651.057	1.29	9.0E+07	Yes
NAC	CL1	GRN	568.134	113.019	569.236	1.42	8.8E+07	Yes
NAC	CL1	UV3	338.284	68.0616	343.136	1.45	7.8E+07	Yes
NAC	CL1	BL2	439.923	29.4692	440.980	1.25	9.0E+07	Yes
NAC	CL1	MT2	727.421	4.11240	727.415	1.34	8.3E+07	Yes
NAC	CL1	CB2	750.505	10.0129	750.495	1.39	7.7E+07	Yes
NAC	CL1	MT3	889.194	10.4720	889.196	--	--	--
NAC	CL1	CB3	937.964	9.54761	937.928	--	--	--
NAC	CL1	MT1	618.945	3.68940	618.949	1.24	9.8E+07	Yes
NAC	CL1	CB1	619.381	9.99526	619.292	1.34	2.9E+03	No
NAC	CL1	CB1A	602.908	9.99526	602.917	--	--	--
NAC	CL1	CB1B	634.531	11.9658	634.526	--	--	--
NAC	CL1	IR3	929.763	66.9995	928.304	1.45	6.2E+07	Yes
NAC	CL1	IR1	751.894	152.929	750.048	1.44	7.1E+07	Yes
NAC	RED	CL2	650.086	149.998	648.879	1.4	8.4E+07	Yes
NAC	RED	GRN	601.032	51.9801	600.959	1.31	9.2E+07	Yes
NAC	RED	MT2	726.633	2.33906	726.624	--	--	--
NAC	RED	CB2	744.255	4.22393	743.912	--	--	--
NAC	RED	MT1	618.911	3.69858	618.922	--	--	--
NAC	RED	CB1	619.568	9.07488	619.481	--	--	--
NAC	RED	IR3	695.435	2.04887	695.040	--	--	--
NAC	RED	IR1	701.900	44.9603	701.692	1.37	7.8E+07	Yes
NAC	BL1	CL2	450.851	102.996	455.471	1.33	1.3E+08	Yes
NAC	BL1	GRN	497.445	5.00811	497.435	1.27	1.1E+07	Yes
NAC	BL1	UV3	386.571	14.0295	389.220	--	--	--
NAC	BL1	BL2	440.035	29.6733	441.077	--	--	--
NAC	UV2	CL2	297.880	59.9535	306.477	--	--	--
NAC	UV2	UV3	315.623	28.9282	317.609	1.34	8.6E+07	Yes
NAC	UV1	CL2	258.098	37.9542	266.321	1.29	8.9E+07	Yes
NAC	UV1	UV3	350.697	9.07263	353.878	--	--	--
NAC	HAL	CL2	655.663	9.26470	655.621	1.25	9.0E+07	Yes
NAC	HAL	GRN	648.028	5.58862	647.808	--	--	--
NAC	HAL	CB1	650.567	2.73589	650.466	--	--	--
NAC	HAL	IR1	663.476	5.25757	663.431	--	--	--
NAC	IR4	CL2	1002.40	35.9966	1001.91	1.53	6.1E+07	Yes
NAC	IR4	IR3	996.723	36.0700	996.460	1.4	6.5E+07	Yes
NAC	IR2	CL2	861.962	97.0431	861.066	1.56	5.9E+07	Yes
NAC	IR2	MT3	889.176	10.4655	889.176	--	--	--
NAC	IR2	CB3	933.657	3.71709	933.593	--	--	--

Camera	Filter 1	Filter 2	Central Wavelength (nm)	Bandpass FWHM (nm)	Effective Wavelength (nm)	PSF FWHM (pixels)	PSF Dynamic Range	Extended PSF?
NAC	IR2	IR3	901.843	44.0356	901.630	1.39	6.7E+07	Yes
NAC	IR2	IR1	827.438	28.0430	827.331	1.34	7.8E+07	Yes
WAC	CL1	CL2	634.928	285.999	633.817	1.72	1.8E+07	Yes
WAC	CL1	RED	648.422	150.025	647.239	1.41	3.4E+07	Yes
WAC	CL1	GRN	567.126	123.999	568.214	1.19	4.9E+07	Yes
WAC	CL1	BL1	460.418	62.2554	462.865	1.46	2.2E+07	Yes
WAC	CL1	VIO	419.684	18.1825	419.822	1.12	2.1E+07	Yes
WAC	CL1	HAL	656.401	9.96150	656.386	1.08	1.4E+08	No
WAC	CL1	IR1	741.456	99.9735	739.826	1.57	4.3E+07	Yes
WAC	IR3	CL2	917.841	45.3074	916.727	1.49	2.6E+07	Yes
WAC	IR3	RED	690.604	3.04414	689.959	--	--	--
WAC	IR3	IR1	790.007	3.02556	783.722	--	--	--
WAC	IR4	CL2	1002.36	25.5330	1001.88	1.48	3.7E+07	No
WAC	IR5	CL2	1034.49	19.4577	1033.87	1.37	4.7E+07	No
WAC	CB3	CL2	938.532	9.95298	938.445	1.77	2.1E+07	No
WAC	MT3	CL2	890.340	10.0116	890.332	1.64	3.6E+07	No
WAC	CB2	CL2	752.364	10.0044	752.354	1.38	3.6E+07	No
WAC	CB2	RED	747.602	4.07656	747.317	--	--	--
WAC	CB2	IR1	752.324	10.0026	752.314	--	--	--
WAC	MT2	CL2	728.452	4.00903	728.418	1.34	1.5E+08	No
WAC	MT2	RED	727.517	2.05059	727.507	--	--	--
WAC	MT2	IR1	728.293	4.00906	728.284	--	--	--
WAC	IR2	CL2	853.258	54.8544	852.448	1.61	6.3E+07	No
WAC	IR2	IR1	826.348	26.0795	826.255	4.67	3.1E+06	No

Table A.2: ISS bandpasses and PSF properties for non-polarized filters. Effective wavelength is the central wavelength of the bandpass convolved with the solar spectrum. Values for CB1A and CB1B in the NAC refer to the two distinct peaks of the CB1 filter response curve. All values were derived from ground-based measurements of the CCD quantum efficiency and filter and optics transmission curves. Adapted and updated from Porco et. al. (2004).

Index	Time		Index	Time		Index	Time		Index	Time
0	0		16	150		32	3200		48	68000
1	5		17	180		33	3800		49	82000
2	10		18	220		34	4600		50	100000
3	15		19	260		35	5600		51	120000
4	20		20	320		36	6800		52	150000
5	25		21	380		37	8200		53	180000
6	30		22	460		38	10000		54	220000
7	35		23	560		39	12000		55	260000
8	40		24	680		40	15000		56	320000
9	50		25	820		41	18000		57	380000
10	60		26	1000		42	22000		58	460000
11	70		27	1200		43	26000		59	560000
12	80		28	1500		44	32000		60	680000
13	90		29	1800		45	38000		61	1000000
14	100		30	2000		46	46000		62	1200000
15	120		31	2600		47	56000		63	No-op

Table A.3: Full list of commandable exposure times in milliseconds.

Gain Index	Description	GAIN_MODE_ID	NAC Gain (e ⁻ /DN)	Ratio (g./g.)	WAC Gain (e-/DN)	Ratio (g./g.)
0	Designed for 4x4	'215 ELECTRONS PER DN'		0.135		0.125
1	Designed for 2x2	'95 ELECTRONS PER DN'		0.310		0.291
2	Used for 1x1	'29 ELECTRONS PER DN'	30.27	1.000	27.68	1.000
3	Used for 1x1	'12 ELECTRONS PER DN'		2.357		2.360

Table A.4: Gain states of the ISS. Actual gain values, derived from in-flight measurements, are given as a ratio relative to gain 2. The highest-gain mode, gain 3, was chosen to match the detector read noise. Updated from Porco et. al. (2004).

		Prepare Cycle					Read-out Cycle			
NAC	Shutter Reset & 200 ms Pad	NAC Filter Wheel	Wait	Pad 325 ms	Wait	Flood & Erase	Exposure	NAC Read-out	Pad 262 ms	Wait
WAC	Shutter Reset & 200 ms Pad	Wait	WAC Filter Wheel	Pad 325 ms	Wait	Flood & Erase	Exposure	Wait	WAC Read-out	Pad 262 ms

Table A.5: One ISS image event, consisting of a prepare and readout cycle. Note that the figure is slightly misleading: WAC readout actually starts at the same time as NAC readout, stops if and when the image buffer fills, and then continues reading out when the WAC readout window starts. From Porco et. al. (2004).

Prepare Index	NAC filter wheel time (sec)	WAC filter wheel time (sec)	Exposure Window (sec)	Total Prepare Time (NAC or WAC)	Total Prepare Time (BOTSIM)
0	1	1	2	4.475	5.475
1	2	2	2	5.475	7.475
2	3	3	2	6.475	9.475
3	5	5	2	8.475	13.475
4	5	5	5	11.475	16.475
5	5	5	13	19.475	24.475
6	5	5	21	27.475	32.475
7	5	5	37	43.475	48.475
8	5	5	53	59.475	64.475
9	5	5	85	91.475	96.475
10	5	5	117	123.475	128.475
11	5	5	181	187.475	192.475
12	5	5	245	251.475	256.475
13	5	5	501	507.475	512.475
14	5	5	101	1019.475	1024.475
15	5	5	1201	1207.475	1212.475

Table A.6: ISS prepare cycle indices and corresponding times.

Readout Index	NAC Readout Window (sec)	WAC Readout Window (sec)	Total Readout Time (NAC or WAC)	Total Readout Time (BOTSIM)
0	50	50	50.525	100.525
1	50	25		75.525
2	50	14		64.525
3	50	6		56.525
4	25	50		75.525
5	25	25	25.525	50.525
6	25	14		39.525
7	25	6		31.525
8	14	50		64.525
9	14	25		39.525
10	14	14	14.525	28.525
11	14	6		20.525
12	6	50		56.525
13	6	25		31.525
14	6	14		20.525
15	6	6	6.525	12.525

*Table A.7: ISS readout cycle indices and corresponding readout times for the maximum data telemetry rate of 48 packets/sec. To calculate the actual readout time, it is necessary to scale by the actual telemetry rate, i.e.: readout time = Integer[readout window * (48 packets/sec TLM rate)/(actual TLM rate)] + pad (525 ms). The relationship between the TELEMETRY_FORMAT_ID from the image label and the corresponding rate in Kbits/sec and packets/sec is given in Table A.8.*

Telemetry Mode (TELEMETRY_FORMAT_ID)	Kbits/sec (INSTRUMENT_DATA_RATE)	Packets/sec
S_N_ER_5	356.6	48
S_N_ER_6	203.1	32
S_N_ER_3	182.8	24
S_N_ER_1	121.9	16
S_N_ER_2	60.9	8

Table A.8: Data telemetry rates used by Cassini. Note that the TELEMETRY_FORMAT_ID keyword in the image labels is invalid for all images taken after SCLK 1431917802 in C37 due to a flight software change. See also the table below.

FLIGHT_SOFTWARE_VERSION_ID	SCLK	Change(s)
1.2	(launch)	<ul style="list-style-type: none"> One overclocked and extended pixel per line Overclocked and extended pixels not valid for LOSSY-compressed images
1.3	1401927444	<ul style="list-style-type: none"> Eight overclocked pixels per line, recorded as two values (averages of the first 2 and last 6, respectively) BIAS_STRIP_MEAN changed appropriately to use new overclocked pixel values Still only one extended pixel per line
1.4	1431917802	<ul style="list-style-type: none"> Extended pixels now a sum of up to 8 pixels in 1x1 mode, 4 in 2x2 and 2 in 4x4 Extended and overclocked pixels for LOSSY-compressed images now returned for last compression block; otherwise set to 0 BIAS_STRIP_MEAN and DARK_STRIP_MEAN for LOSSY images now derived from data returned in last compression block FILTER_TEMPERATURE and SENSOR_HEAD_ELECTRONICS_TEMPERATURE no longer valid (now recorded at time of downlink, not exposure) TELEMETRY_FORMAT_ID no longer valid

Table A.9: Flight software changes since launch. "SCLK" is the spacecraft clock count at the time of the change.




Figure A. 1: Full-system transmission for medium and broadband ISS filters. For the NAC: purple corresponds to UV1 (solid), UV2 (dotted), and UV3 (dashed), blue to BL1 (solid) and BL2 (dashed), green to GRN, red to RED, and orange to IR1 (solid), IR2 (dot-dashed), IR3 (dashed) and IR4 (dotted). For the WAC: purple for VIO, blue for BL1, green for GRN, red for RED, and orange for IR1 (solid), IR2 (dashed) and IR3 (dotted); IR4 and IR5 not shown. From Porco et. al. (2004).




Figure A. 2: System transmission for ISS narrow-band filters. For the NAC: purple for MT1 (solid), MT2 (dotted) and MT3 (dashed), green for CB1 (solid), CB2 (dotted) and CB3 (dashed), and red for HAL. For the WAC: purple for MT2 (solid) and MT3 (dashed), green for CB2 (solid) and CB3 (dashed), and red for HAL. The geometric albedo of Titan (black) is also shown to illustrate the placement of filters relative to atmospheric methane features. From Porco et. al. (2004).




Figure A. 3: System transmissions for select combined filters. For the NAC: purple for UV2/UV3, green for RED/GRN and red for RED/IR1 (solid), IR1/IR2 (dotted), IR2/IR3 (dashed) and IR3/IR4 (dot-dashed). From Porco et. al. (2004).




Figure A. 4: Transmission of ISS polarizing filters: visible-wavelength (NAC P0, P60 and P120, solid lines) and infrared (NAC IRP0 and WAC IRP0 and IRP90, dashed). Transmission is shown for both perpendicular and parallel orientations of the polarization axis to that of the incoming light. The IR perpendicular case is not shown; its transmission peaks at 0.00225 and is under 0.001 for wavelengths longward of 630 nm.




Figure A. 5: ISS optics transmission. From Porco et. al. (2004).




Figure A. 6: ISS CCD quantum efficiency (electrons per incident photon) as a function of wavelength. From Porco et. al. (2004).

Appendix B: ISS Observation Descriptions

As explained in Section 3, each ISS observation is given a unique identifier, the observation request name, which is recorded in the image label as the OBSERVATION_ID and has the format:

$$[INST]_[REV][TI]_[UNIQUENAME]_[PINST]$$

where *INST* is the abbreviation for the instrument being commanded (always “ISS,” in our case), *REV* is the rev number, *TI* is a 2-character target identifier (see Table 3 in Section 3), *UNIQUENAME* is a unique descriptor chosen by the observation creator, and *PINST* indicates the name of the prime instrument if other than the ISS.

The *UNIQUENAME* portion of the OBSERVATION_ID is left up to the observation creator and thus follows no fixed convention, but it can often be quite descriptive, particularly for certain types of observations that are repeated many times over the course of the mission. The following is a partial list of these observations, organized by target discipline, which together constitute the majority of ISS data. In most cases, the bold-faced names shown are not the complete *UNIQUENAME*, but rather a sub-string contained within it, which can be used for querying the OBSERVATION_ID keyword. Asterisks indicate a wildcard character which matches an arbitrary character string.

Saturn Atmosphere

1X2WPH25: A 1X2 WAC mosaic of Saturn using narrow and medium-band filters for methane-band photometry and color and polarization. The intent is to gather images over time to accumulate phase angle coverage and to record seasonal changes.

APPRMOV: Saturn/rings approach movie. Clear and methane filters are used to get vertical cloud structure. Broadband filters are sometimes used to see cloud color. The intent is to gather images over time to see cloud motion.

[**ATMDYN**, **COMPSIT**, **FTRACK**, **GLOBDYN**, **GLOMAP**, **HEMDYN**, **HRES**, **MAP**, **NHEMMOVIE**, **POLECAM**, **POLEDYN**, **POLRMOV**]: Ride-along observations with other instruments (VIMS, UVIS, CIRS). The pointing is determined by the other instrument. Depending on

the pointing, the observations are intended to observe night side in search for lightning and aurora, or the day side for cloud structure and motion.

AUR: Night side observations, some of them in different filters. The intent is to observe lightning and aurora changing in time. Some of those observations are ride-alongs with other instruments (UVIS, VIMS), for which the pointing is defined by the other instrument. Depending on the pointing, the observations may, instead of the night side, be on the day side and observe cloud structure and motion.

MONIT: Frequent Saturn observations at low spatial resolution with different filters used for color and vertical cloud structure. The intent is to search for newly appearing thunderstorms and comet impacts.

NALGTNG: NAC lightning/aurora search at the night side of Saturn. Some of those observations use different sets of filters to get spectrum of lightning and aurora.

WALGTNG: WAC lightning/aurora search at the night side of Saturn. Some of those observations use different sets of filters to get spectrum of lightning and aurora.

WIND: Part of Saturn wind measurement campaign. Simultaneous observations are coordinated between ISS, VIMS, and CIRS. ISS observes the same area on Saturn on consecutive rotations at high spatial resolution with NAC. The intent is to see cloud motion at a 10-hour timescale, and to cover all latitudes by the set of such observations.

Saturn Rings

ARCLELR: Observations of faint ring arcs (Anthe, Methone, Pallene, and/or Aegaeon) at low elevation angle and low spatial resolution

ARCORBIT: Observations of faint ring arcs (Anthe, Methone, Pallene, and/or Aegaeon) to determine their orbits

AZSCAN: Azimuthal scan around the rings. Early scans used denser spacing than later scans. Azscans were discontinued during the XM due to wear on the reaction wheels.

DIFFRING: Observations of diffuse rings

DIFFUSRNG: Search for diffuse rings

DIFRNGAPR: Search for diffuse rings

DIFSATSRC: Search for diffuse rings and satellites

DRCLOSE: Observations of the D ring at high spatial resolution

DRING: Observe the D ring

DRLPMOV: Movie of the D ring at low phase angle

DRNGMOV: Movie of the D ring

DUSTHAZRD: Observations of the G ring while spacecraft was oriented so as to avoid dust hazards

EGAPMOVMP: Movie of the Encke Gap at moderate phase angle

EGPHASE: Photometric observations of the E ring and/or G ring, usually at low elevation angle

EGxxPHASE: Observe the E ring and/or the G ring at xx degrees phase angle

ENCKEMOV: Movie of the Encke Gap

ERNGVERT: Edge-on observations of the E ring

ERNGVRLP: Edge-on observations of the E ring at low phase angle

ExxPHASE: Observe the E ring at xx degrees phase angle

FNTLPLE: Observations of faint rings at low phase angle and low elevation angle

FNTLPMOV: Movie of faint rings at low phase angle

FxxPHASE: Observe the F ring at xx degrees phase angle

GARCORB: Track the orbit of the arc within the G ring

GARCPROM: Unrelated observations of Prometheus and of the arc within the G ring (including Aegaeon) combined into one maneuver

GRINGARC: Track the orbit of the arc within the G ring

GxxPHASE: Observe the G ring at xx degrees phase angle

HIPHAASE: Observations at high phase angle

HIPHNAC: NAC observations at high phase angle

HIPHWAC: WAC observations at high phase angle

HPLELR: Observations of diffuse rings at high phase angle, low elevation angle, and low spatial resolution

HPMELR: Observations of diffuse rings at high phase angle, moderate elevation angle, and low spatial resolution

LOWPHASE: Observations at low phase angle

LRHPENKMOV: Movie of the Encke Gap at low spatial resolution and high phase angle

LRLEMP: Observations of diffuse rings at low spatial resolution, low elevation angle, and moderate phase angle

*MOV: Ring ansa movie. Examples include, BMOVIE, SHRTMOV. Point-and-stare at one ring ansa while acquiring images.

PHOEBEHOL: Observations of Saturn's shadow cast onto the Phoebe ring

PROMPROP: Unrelated observations of Prometheus and of "propeller" moons in the A ring combined into one maneuver

PROPELLR: Observe "propeller" moons in the A ring

PROPRETRG: Re-targetting known "propeller" moons to track their orbits

PROPSURVY: Azimuthally-complete (or nearly so) survey of the A ring for "propeller" moons

RHEARXPXP: Observations to search for material orbiting Rhea during ring-plane crossing at low phase

RHEARPXHP: Observations to search for material orbiting Rhea during ring-plane crossing at high phase

SHADBNDLE: Observe the boundary of Saturn's shadow cast onto the rings at low elevation angle

SHADBOUND: Observe the boundary of Saturn's shadow cast onto the rings

SPK: Basically, any observation that starts with "SPK" is a spoke obs. They tend to have suffixes appended to them. Common ones are DF/LF for Dark Face/Lit Face, HP/LP for High Phase/Low Phase, and LR/MR/HR for Low Resolution/ Medium Resolution/High Resolution.

SPKMOV: Spoke periodicity movies. These are 'point and stare' observations where we select a ring ansa and image at regular intervals (usually with the WAC) to watch spoke move through the field of view.

SPKTRK: Spoke tracking movie. Instead of a point and stare, this type of spoke obs tracks the motion of the spokes.

SPKFORM, SPK*FORM: Spoke formation movie. Like a SPKTRK, but with extremely rapid imaging using 4x4 summation. These were designed to catch a spoke forming, which was thought to be very fast, based on certain Voyager observations. Discontinued in the XXM.

URBETORI: Observations of the main rings during a UVIS-led occultation of the star Beta Orionis

Titan

CLOUD: These observations are typically used during non-targeted encounters or in the 2-3 days preceding or following a Titan targeted flyby. These images are usually taken at a distance of 400,000 to 1 million kilometers. They are designed to fill gaps in our map of Titan and to monitor clouds in Titan's troposphere for periods up to 24 hours. Depending on the geometry of the observation, this type of sequence may use the name SPOLE or NPOLE (or even just TITAN, as was used during Rev088). HIGHRESNA (or HIGHRES), REGMAPNA (or REGMAP), GLOBMAPNA (or GLOBMAP), and

MONITORNA (or MONITOR): Mosaic sequences designed to map Titan surface and detect small-scale cloud features during a Titan targeted encounter. The name used pertains to the distance from Titan the observation is acquired from, with HIGHRES acquired at close range (typically less than 100,000 kilometers) and MONITOR taken the furthest out, with nine, narrow-angle-camera frames covering the entire visible disk (on a typical usage, like Ta). For each footprint in these mosaics, several CL1-CB3 images are taken, so they can be summed on the ground to increase the signal-to-noise ratio, along with one CL1-CB3 image.

MxxxRxHZxxx: This is a typical observation name for the ISS Titan haze monitoring campaign. These observation names are a mashup of three parts: Mxxx represents the approximate phase angle of the observation at 30 degree increments, Rx represents the range and thus what instruments ride along with ISS (R1 being the closest and R3 being the farthest), and the final three numbers represent the day-of-year during which the observation was taken. Example: "M120R2HZ111." Use HZ in the OBSERVATION_ID keyword to find these observations.

MxxxRxCLDxxx: This is a typical observation name for the ISS Titan cloud monitoring campaign. These observation names are a mashup of three parts: Mxxx represents the approximate phase angle of the observation at 30 degree increments, Rx represents the range and thus what instruments ride along with ISS (R1 being the closest and R3 being the farthest), and the final three numbers represent the day-of-year during which the observation was taken. Example: "M60R1CLD317." Use CLD in the OBSERVATION_ID keyword to find these observations.

PHOTOMWAC: A single WAC frame of Titan using narrow and medium-band filters for methane-band photometry and color and polarization. The intent is to gather images over time to accumulate phase angle coverage and to record seasonal changes.

TEA: This sequence has only been used during the Cassini Solstice Mission. These are longer sequences designed to monitor cloud features on Titan over a several day period.

Encounter Ride-along Observations: ISS typically acquires images while UVIS, CIRS, or VIMS are prime during Titan encounters. These are often taken while Cassini is over Titan's night side and ISS studies the moon's upper haze layers. A larger span of filter combinations is used. CIRS ride-along observations have names like FIRNADMAP, MIRLMBINT, FIRNADCMP, or MIDIRTMAP. VIMS ride-along observations have names like HIRES, GLOBMAP, and CLOUDMAP (and like ISS prime surface observations, are taken while over Titan's dayside). UVIS ride-along observations use names like EUVFUV or HDACSTARE.

Icy Satellites

GEOLOG: A high-resolution image or mosaic that is intended for the study of detailed surface geology of an icy surface or a particular terrain on it. These images generally use 1x1 NAC CL1-CL2 images or mosaics and may include a 1x1 WAC CL1-CL2 image for spatial reference.

GLOCOL: A multi-color series of NAC 1x1 satellite images, each often showing the whole-disk within a NAC frame, that are intended to provide broad multi-spectral global color coverage. A standard choice of filters for a general global color series would include 24 NAC colors, 9 NAC polarizers, 4 WAC polarizers.

HILLSPHE: A wide-field mosaic surrounding a satellite intended to search for small objects in its Hill Sphere.

HIRES: A solitary image or collection of images (sometimes a mosaic) obtained at especially high spatial resolution (often a couple tens of meters per pixel or better). When a “HIRES” mosaic is obtained, it may also be part of a satellite close flyby sequence for which a lower resolution “LOWRES” mosaic was also obtained that shows the same or adjacent geographic region from a different viewing geometry. Often, the LOWRES counterpart includes images obtained using a wider range of color filters than could be obtained over the short duration of the highest-resolution coverage opportunity.

LOWPHASE: A series of successive images intended to show how the photometric and polarization phase curves change near opposition (i.e. at relatively small phase angles). The observations generally show how the brightness of an icy object changes at phase angles less than 10-degrees, but can include the transition to larger phase angles up to a few tens of degrees. The series usually allows for multispectral imaging and sometimes GRN polarization frames that are repeated as the phase angle incrementally changes.

LOWRES: A mosaic or series of color mosaics obtained at the approach or receding phase of a close-flyby. The low resolution mosaic (usually tens to hundreds of meters/pixel) may show the same or an adjacent region to a corresponding high-resolution mosaic (see HIRES) that is obtained at closest-approach. The longer duration over which a low-resolution mosaic may be obtained allows more time for the use of multi-spectral and polarization filters than the brief HIRES series where often only 1x1 NAC CL1-CL1 images can be acquired.

MORPH: A high-resolution image, mosaic, or series of images intended to investigate the detailed morphology of a specific geological feature or terrain. These are almost always 1x1 NAC CL1/CL2 images, but sometimes include a 1x1 WAC CL1/CL2 for spatial reference.

MUTUALEVE: MUTUAL EVENT observation. A movie series of 1x1 NAC images (usually CL1/CL2) that incrementally shows the passing of two satellite bodies through the camera field of view. The main purpose was to refine satellite ephemerides. However, the images are often well-exposed enough to be useful for photometric studies.

OT_L5: A series of images to search for possible new satellites residing in various Lagrange points.

OT_OUTERSATS: Designation for a series of images that were used to obtain basic photometric and color information for small outer satellites of Saturn.

PHOTOPOL: A series of 1x1 LOSSLESSLY COMPRESSED 8- or 12-bit NAC images that were obtained for the purpose of investigating the PHOTometric and POLARization behavior of icy satellite surfaces. The minimal filter choice for polarization was P0,P60,P120 filters used in combination with GRN, and sometimes UV3 and MT2 as allowed by time and data volume. For photometry, the minimum filter combination was generally CL1-CL2, -UV3,-GRN,-IR3.

REGMAP: Usually a multi-image mosaic intended to obtain REGIONAL MAPPING coverage at moderate to high-resolution (typically on scales of tens to hundreds of pixels/kilometer). The mosaic series minimally contains a multi-panel layout of NAC 1x1 CL1-CL2 images, and often extends to other NAC wavelengths as was allowed by available data volume and time limitations. One or more CL1-CL2 WAC images was often included for context and mosaic panel co-registration.

ROTCOL: A series of images intended to observe how the color, albedo, and shape of an irregularly-shaped satellite, like Hyperion, changes with longitude and latitude. The filter choices generally consist of NAC CL1-CL2, CL1-UV3, CL1-GRN, CL1-IR3, and sometimes a GRN NAC polarization set, all repeated over time as the object rotates and the spacecraft changes position.

SATELLORB: A movie series of NAC 1x1 CL1-CL2 images obtained for the purpose of refining satellite ephemerides. The series usually shows a small satellite moving over a substantial portion of its orbit around Saturn.

STEREO: Generally a pair of high-resolution mosaics of an icy satellite surface intended to be used for stereophotogrammetry and the creation of digital elevation maps to study topography. These were almost always made up of NAC 1x1 CL1-CL2 images.

xxxWyyyPH (example 094W105PH): A series of whole-disk observations designed to obtain multispectral photometric and polarization coverage of each satellite over all longitudes and phase angles. The three 'xxx' digits identify the central longitude and the 'yyy' digits identify the phase angle. The aim was to obtain complete coverage in longitude increments of 60-degrees and phase angle increments of 15-degrees over as wide a range of phase angle as possible. The standard filter set was 4 NAC color + 9 NAC polarizer + 4 WAC polarizer filters + and three CL1-CL2 NAC images, one at the beginning, one in the middle, and one at the end.

xxxxxCA (for example, ENCELCA): A sequence of single NAC CL1-CL2 filter images obtained in rapid succession during the closest-approach portion of a satellite close flyby. Due to the rapidly changing scene geometry during a flyby, image formats and exposure times were optimized to avoid smear and allow as many frames to be targeting in a short time as possible. Thus, depending on the time and data volume resources available, the images can be 1x1, 8-bit losslessly compressed images, 1x1 8-bit lossy-compressed images, and occasionally 2x2 lossless or lossy images. In some cases, NAC CL1-UV3, CL1-GRN, and CL1-IR3 frames were obtained.

ZEROPHASE: A series of incremental images that show how the brightness of an icy object changes at the phase angle phases through zero degrees phase, that is, through "opposition". The observation is intended to measure the opposition effect for each satellite – a non-linear surge in brightness with decreasing phase angle to zero that is observed on regolith covered planetary surfaces.

Small Satellites (aka “Rocks”)

FMOVIE: A stare in a fixed orientation to image F ring material passing through the NAC field of view.

FRSTRCHAN: Prometheus & F ring streamer-channel movie. NAC tracks Prometheus to follow the evolution of the streamer-channel feature raised in the F ring by gravitational interaction with Prometheus at the satellite's apoapse passage.

MUTUALEVE: A satellite mutual event observation. NAC images of one satellite transiting another. Exposure times are chosen for satellite surfaces.

SATELORB: Astrometry observations of the smaller saturnian satellites or “rocks” for orbit determination and monitoring purposes. NAC targets specified satellite, multiple satellites targeted in each observation. NAC images always use CL1/CL2 filters and 680 ms exposures. (Exposure times are chosen for the star background, not for the satellite surface.) Satellite is usually/often saturated, though this doesn't adversely affect the image for astrometry purposes. A second NAC is usually taken with a shorter exposure to image the satellite without saturating.

Enceladus Plume

PLM*: Standard plume observation, taking plume images in a range of filters. Phase angle is typically between 130 and 165 degrees. These observations are for measuring the plume phase function as well as tracking the individual jets back to their sources when the resolution is sufficient.

PLMPHS*: Plume phase angle campaign. These observations are part of a series where the geometry was nearly the same, except for the phase angle. The purpose is again to measure the plume phase function.

PLMSCECL: Plume observation while the spacecraft is in eclipse. These observations allow for phase angles higher than 165 deg, as the Sun is obscured by Saturn.

Calibration

CALSTAR: A calibration observation of photometric standard stars. CALSTAR1 is for Vega, CALSTAR2 is 77/78 Tau, and CALSTAR3 is HR 996. This is a 2X2 mosaic centered on the star such that it falls on four different positions on the detector in order to improve sampling.

IC_DARKCAL: A calibration observation of images taken with the shutter closed to measure dark signal, taken over a range of exposure times.

Appendix C: Rev/Orbit and Sequence Boundaries

Rev/Orbit Boundaries:

Rev	Start Time (UTC)	End Time (UTC)
C27	2001-189T23:00:00	2001-252T23:00:00
C28	2001-251T23:00:00	2001-309T00:00:00
C29	2001-309T00:00:00	2002-014T00:00:00
C30	2002-014T00:00:00	2002-070T00:00:00
C31	2002-070T00:00:00	2002-125T23:00:00
C32	2002-125T23:00:00	2002-188T23:00:00
C33	2002-188T23:00:00	2002-265T21:40:00
C34	2002-265T21:40:00	2002-334T11:36:00
C35	2002-334T11:36:00	2003-039T08:03:00
C36	2003-039T08:03:00	2003-119T03:16:00
C37	2003-119T03:16:00	2003-166T01:17:00
C38	2003-166T00:52:00	2003-222T19:32:00
C39	2003-222T19:32:00	2003-292T15:05:00
C40	2003-292T15:05:00	2004-009T12:25:00
C42	2004-009T12:25:00	2004-051T00:27:00
C41	2004-010T00:00:00	2004-020T00:00:00
C43	2004-051T00:27:00	2004-092T21:28:00
C44	2004-092T21:28:00	2004-135T17:40:00
0	2004-135T17:40:00	2004-240T07:56:00
00A	2004-240T07:56:00	2004-326T08:41:00
1	2004-257T12:04:15	2004-358T08:10:44
00B	2004-326T08:41:00	2004-366T07:02:00
2	2004-358T08:10:44	2005-032T03:41:35
CIA	2004-366T00:00:00	2005-002T01:00:00

Rev	Start Time (UTC)	End Time (UTC)
00C	2004-366T07:02:00	2005-032T03:25:00
CRT	2005-006T11:53:00	2005-015T17:17:00
BGC	2005-017T12:30:00	2005-022T10:38:00
3	2005-032T03:25:00	2005-058T06:19:00
4	2005-058T06:19:00	2005-078T17:34:00
5	2005-078T17:34:00	2005-096T22:32:00
6	2005-096T22:32:00	2005-113T22:32:00
7	2005-113T22:32:00	2005-132T02:52:00
8	2005-132T02:52:00	2005-150T07:21:00
9	2005-150T07:21:00	2005-168T12:12:00
10	2005-168T12:12:00	2005-186T17:57:00
11	2005-186T17:57:00	2005-205T00:58:00
12	2005-205T00:58:00	2005-223T07:30:00
13b	2005-212T21:00:00	2005-242T20:43:00
13	2005-223T07:30:00	2005-240T11:01:00
14	2005-240T11:01:00	2005-257T15:49:00
15	2005-257T15:49:00	2005-275T22:33:00
16	2005-275T22:33:00	2005-293T22:59:00
17	2005-293T22:59:00	2005-317T04:57:00
18	2005-317T04:57:00	2005-345T04:26:00
19	2005-345T04:26:00	2006-005T14:08:00
20	2006-005T14:08:00	2006-036T20:58:00
21	2006-036T20:58:00	2006-068T03:34:00
22	2006-068T03:34:00	2006-099T09:03:00

Rev	Start Time (UTC)	End Time (UTC)
23	2006-099T09:03:00	2006-130T15:31:00
24	2006-130T15:31:00	2006-161T22:06:00
25	2006-161T22:06:00	2006-193T04:18:00
26	2006-193T04:18:00	2006-216T20:25:00
27	2006-216T20:25:00	2006-240T18:13:00
28	2006-240T18:13:00	2006-260T16:06:00
29	2006-260T16:06:00	2006-276T17:53:00
24a	2006-263T19:22:00	2006-265T19:07:00
24b	2006-265T19:07:00	2006-295T17:26:00
30	2006-276T17:53:00	2006-292T21:09:00
31	2006-292T21:09:00	2006-307T00:09:00
32	2006-307T00:09:00	2006-318T23:34:00
33	2006-318T23:34:00	2006-330T22:23:00
34	2006-330T22:23:00	2006-342T21:07:00
35	2006-342T21:07:00	2006-356T23:27:00
36	2006-356T23:27:00	2007-008T04:37:00
37	2007-008T04:37:00	2007-024T12:27:00
38	2007-024T12:27:00	2007-041T10:51:00
39	2007-041T10:51:00	2007-058T07:48:00
40	2007-058T07:48:00	2007-074T11:41:00
41	2007-074T11:41:00	2007-090T16:09:00
42	2007-090T16:09:00	2007-106T19:51:00
43	2007-106T19:51:00	2007-122T22:33:00
44	2007-122T22:33:00	2007-139T00:10:00
45	2007-139T00:10:00	2007-155T00:20:00
46	2007-155T00:20:00	2007-171T00:17:00
47	2007-171T00:17:00	2007-190T08:51:00
48	2007-190T08:51:00	2007-221T15:31:00

Rev	Start Time (UTC)	End Time (UTC)
49	2007-221T15:31:00	2007-257T10:46:00
50	2007-257T10:46:00	2007-285T09:12:00
51	2007-285T09:12:00	2007-309T08:14:00
52	2007-309T08:14:00	2007-329T07:59:00
53	2007-329T07:59:00	2007-345T05:32:00
54	2007-345T05:32:00	2007-361T01:24:00
55	2007-361T01:24:00	2008-009T22:29:00
56	2008-009T22:29:00	2008-021T21:12:00
57	2008-021T21:12:00	2008-033T20:00:00
58	2008-033T20:00:00	2008-045T19:27:00
59	2008-045T19:27:00	2008-056T22:09:00
60	2008-056T22:09:00	2008-067T13:03:00
61	2008-067T13:03:00	2008-078T03:38:00
62	2008-078T03:38:00	2008-088T00:28:00
63	2008-088T00:28:00	2008-097T14:24:00
64	2008-097T14:24:00	2008-107T04:06:00
65	2008-107T04:06:00	2008-116T17:22:00
66	2008-116T17:22:00	2008-126T06:35:00
67	2008-126T06:35:00	2008-134T23:09:00
68	2008-134T23:09:00	2008-142T22:21:00
69	2008-142T22:21:00	2008-150T08:02:00
70	2008-150T08:02:00	2008-157T10:46:00
71	2008-157T10:46:00	2008-164T13:29:00
72	2008-164T13:29:00	2008-171T16:02:00
73	2008-171T06:02:00	2008-178T18:35:00
74	2008-178T18:35:00	2008-185T21:03:00
75	2008-185T21:03:00	2008-192T22:20:00
76	2008-192T22:20:00	2008-199T22:27:00

Rev	Start Time (UTC)	End Time (UTC)
77	2008-199T22:27:00	2008-206T22:40:00
78	2008-206T22:40:00	2008-213T18:50:00
79	2008-213T18:50:00	2008-221T04:07:00
80	2008-221T04:07:00	2008-228T13:08:00
81	2008-228T13:08:00	2008-235T21:42:00
82	2008-235T21:42:00	2008-243T06:40:00
83	2008-243T06:40:00	2008-250T15:33:00
84	2008-250T15:33:00	2008-258T00:24:00
85	2008-258T00:24:00	2008-265T09:14:00
86	2008-265T09:14:00	2008-272T18:03:00
87	2008-272T18:03:00	2008-280T02:35:00
88	2008-280T02:35:00	2008-287T10:41:00
89	2008-287T10:41:00	2008-294T18:07:00
90	2008-294T18:07:00	2008-302T01:15:00
91	2008-302T01:15:00	2008-309T23:29:00
92	2008-309T23:29:00	2008-317T22:29:00
93	2008-317T22:29:00	2008-325T10:17:00
94	2008-325T10:17:00	2008-333T09:31:00
95	2008-333T09:31:00	2008-340T18:21:00
96	2008-340T18:21:00	2008-348T17:48:00
97	2008-348T17:48:00	2008-356T13:04:00
98	2008-356T13:04:00	2008-366T00:11:00
99	2008-366T00:11:00	2009-009T13:58:00
100	2009-009T13:58:00	2009-019T03:43:00
101	2009-019T03:43:00	2009-028T17:09:00
102	2009-028T17:09:00	2009-039T02:36:00
103	2009-039T02:36:00	2009-051T01:22:00
104	2009-051T01:22:00	2009-063T00:20:00

Rev	Start Time (UTC)	End Time (UTC)
105	2009-063T00:20:00	2009-074T22:18:00
106	2009-074T22:18:00	2009-086T03:51:00
107	2009-086T03:51:00	2009-094T00:29:00
108	2009-094T00:29:00	2009-106T22:05:00
109	2009-106T22:05:00	2009-122T00:37:00
110	2009-122T00:37:00	2009-137T12:44:00
111	2009-137T12:44:00	2009-152T23:44:00
112	2009-152T23:44:00	2009-168T13:08:00
113	2009-168T13:08:00	2009-184T04:04:00
114	2009-184T04:04:00	2009-199T20:28:00
115	2009-199T20:28:00	2009-215T14:23:00
116	2009-215T14:23:00	2009-231T10:08:00
117	2009-231T10:08:00	2009-251T09:15:00
118	2009-251T09:15:00	2009-275T07:03:00
119	2009-275T07:03:00	2009-296T15:59:00
120	2009-296T15:59:00	2009-315T16:57:00
121	2009-315T16:57:00	2009-334T17:36:00
122	2009-334T17:36:00	2009-352T05:54:00
123	2009-352T05:54:00	2010-003T03:49:00
124	2010-003T03:49:00	2010-019T04:17:00
125	2010-019T04:17:00	2010-035T22:18:00
126	2010-035T22:18:00	2010-053T11:27:00
127	2010-053T11:27:00	2010-071T01:42:00
128	2010-071T01:42:00	2010-088T16:01:00
129	2010-088T16:01:00	2010-107T16:20:00
130	2010-107T16:20:00	2010-128T02:04:00
131	2010-128T02:04:00	2010-146T07:04:00
132	2010-146T07:04:00	2010-162T06:27:00

Rev	Start Time (UTC)	End Time (UTC)
133	2010-162T06:27:00	2010-178T04:23:00
134	2010-178T04:23:00	2010-196T03:21:00
135	2010-196T03:21:00	2010-216T01:12:00
136	2010-216T01:12:00	2010-236T00:18:00
137	2010-236T00:18:00	2010-255T23:44:00
138	2010-255T23:44:00	2010-277T20:02:00
139	2010-277T20:02:00	2010-301T17:21:00
140	2010-301T17:21:00	2010-324T01:56:00
141	2010-324T01:56:00	2010-344T16:15:00
142	2010-344T16:15:00	2010-365T07:30:00
143	2010-365T07:30:00	2011-020T21:28:00
144	2011-020T21:28:00	2011-041T08:05:00
145	2011-041T08:05:00	2011-065T12:39:00
146	2011-065T12:39:00	2011-093T09:51:00
147	2011-093T09:51:00	2011-119T02:06:00
148	2011-119T02:06:00	2011-150T09:15:00
149	2011-150T09:15:00	2011-180T18:56:00
150	2011-180T18:56:00	2011-202T11:10:00
151	2011-202T11:10:00	2011-224T05:01:00
152	2011-224T05:01:00	2011-246T00:47:00
153	2011-246T00:47:00	2011-265T18:58:00
154	2011-265T18:58:00	2011-283T13:40:00
155	2011-283T13:40:00	2011-301T09:10:00
156	2011-301T09:10:00	2011-319T06:12:00
157	2011-319T06:12:00	2011-337T03:13:00
158	2011-337T03:13:00	2011-357T19:14:00
159	2011-357T19:14:00	2012-016T15:46:00
160	2012-016T15:46:00	2012-040T12:16:00

Rev	Start Time (UTC)	End Time (UTC)
161	2012-040T12:16:00	2012-061T03:51:00
162	2012-061T03:51:00	2012-078T22:54:00
163	2012-078T22:54:00	2012-096T18:17:00
164	2012-096T18:17:00	2012-114T13:46:00
165	2012-114T13:46:00	2012-132T08:28:00
166	2012-132T08:28:00	2012-149T04:15:00
167	2012-149T04:15:00	2012-169T00:46:00
168	2012-169T00:46:00	2012-192T22:45:00
169	2012-192T22:45:00	2012-215T09:52:00
170	2012-215T09:52:00	2012-236T15:38:00
171	2012-236T15:38:00	2012-257T22:17:00
172	2012-257T22:17:00	2012-280T09:18:00
173	2012-280T09:18:00	2012-304T07:16:00
174	2012-304T07:16:00	2012-324T02:54:00
175	2012-324T02:54:00	2012-338T10:58:00
176	2012-338T10:58:00	2012-351T17:22:00
177	2012-351T17:22:00	2012-364T23:50:00
178	2012-364T23:50:00	2013-012T06:43:00
179	2013-012T06:43:00	2013-025T13:55:00
180	2013-025T13:55:00	2013-038T21:11:00
181	2013-038T21:11:00	2013-051T05:02:00
182	2013-051T05:02:00	2013-063T04:01:00
183	2013-063T04:01:00	2013-075T02:13:00
184	2013-075T02:13:00	2013-087T01:16:00
185	2013-087T01:16:00	2013-097T16:49:00
186	2013-097T16:49:00	2013-107T06:25:00
187	2013-107T06:25:00	2013-116T20:08:00
188	2013-116T20:08:00	2013-126T09:48:00

Rev	Start Time (UTC)	End Time (UTC)
189	2013-126T09:48:00	2013-135T23:29:00
190	2013-135T23:29:00	2013-146T06:37:00
191	2013-146T06:37:00	2013-158T05:29:00
192	2013-158T05:29:00	2013-170T04:34:00
193	2013-170T04:34:00	2013-182T03:38:00
194	2013-182T03:38:00	2013-196T05:14:00
195	2013-196T05:14:00	2013-217T15:37:00
196	2013-217T15:37:00	2013-241T13:48:00
197	2013-241T13:48:00	2013-270T14:02:00
198	2013-270T14:02:00	2013-311T06:31:00
199	2013-311T06:31:00	2013-351T21:24:00
200	2013-351T21:24:00	2014-019T10:10:00
201	2014-019T10:10:00	2014-051T16:45:00
202	2014-051T16:45:00	2014-083T17:01:00
203	2014-083T17:01:00	2014-117T12:47:00
204	2014-117T12:47:00	2014-151T07:36:00
205	2014-151T07:36:00	2014-183T05:52:00
206	2014-183T05:52:00	2014-215T05:56:00
207	2014-215T05:56:00	2014-247T00:17:00
208	2014-247T00:17:00	2014-278T21:43:00
209	2014-278T21:43:00	2014-318T23:03:00
210	2014-318T23:03:00	2014-358T19:16:00
211	2014-358T19:16:00	2015-025T18:41:00
212	2015-025T18:41:00	2015-057T17:14:00
213	2015-057T17:14:00	2015-087T15:40:00
214	2015-087T15:40:00	2015-115T16:46:00
215	2015-115T16:46:00	2015-139T03:31:00
216	2015-139T03:31:00	2015-158T00:48:00

Rev	Start Time (UTC)	End Time (UTC)
217	2015-158T00:48:00	2015-176T22:56:00
218	2015-176T22:56:00	2015-197T07:19:00
219	2015-197T07:19:00	2015-219T02:41:00
220	2015-219T02:41:00	2015-240T22:28:00
221	2015-240T22:28:00	2015-262T18:28:00
222	2015-262T18:28:00	2015-280T14:12:00
223	2015-280T14:12:00	2015-294T12:01:00
224	2015-294T12:01:00	2015-308T10:42:00
225	2015-308T10:42:00	2015-321T18:22:00
226	2015-321T18:22:00	2015-334T11:52:00
227	2015-334T11:52:00	2015-347T05:22:00
228	2015-347T05:22:00	2015-360T00:24:00
229	2015-360T00:24:00	2016-007T20:48:00
230	2016-007T20:48:00	2016-022T05:37:00
231	2016-022T05:37:00	2016-038T03:23:00
232	2016-038T03:23:00	2016-058T00:08:00
233	2016-058T00:08:00	2016-081T21:04:00
234	2016-081T21:04:00	2016-109T17:23:00
235	2016-109T17:23:00	2016-141T12:31:00
236	2016-141T12:31:00	2016-169T05:13:00
237	2016-169T05:13:00	2016-193T03:17:00
238	2016-193T03:17:00	2016-212T19:05:00
239	2016-212T19:05:00	2016-226T14:01:00
240	2016-226T14:01:00	2016-238T12:51:00
241	2016-238T12:51:00	2016-250T11:51:00
242	2016-250T11:51:00	2016-262T11:01:00
243	2016-262T11:01:00	2016-273T04:54:00
244	2016-273T04:54:00	2016-282T18:29:00

Rev	Start Time (UTC)	End Time (UTC)
245	2016-282T18:29:00	2016-292T08:10:00
246	2016-292T08:10:00	2016-301T21:45:00
247	2016-301T21:45:00	2016-311T12:16:00
248	2016-311T12:16:00	2016-320T09:10:00
249	2016-320T09:10:00	2016-328T08:21:00
250	2016-328T08:21:00	2016-335T23:32:00
251	2016-335T23:32:00	2016-343T03:45:00
252	2016-343T03:45:00	2016-350T07:57:00
253	2016-350T07:57:00	2016-357T11:57:00
254	2016-357T11:57:00	2016-364T15:56:00
255	2016-364T15:56:00	2017-005T19:49:00
256	2017-005T19:49:00	2017-012T23:42:00
257	2017-012T23:42:00	2017-020T03:34:00
258	2017-020T03:34:00	2017-027T07:29:00
259	2017-027T07:29:00	2017-034T11:38:00
260	2017-034T11:38:00	2017-041T16:06:00
261	2017-041T16:06:00	2017-048T20:28:00
262	2017-048T20:28:00	2017-056T00:31:00
263	2017-056T00:31:00	2017-063T04:35:00
264	2017-063T04:35:00	2017-070T08:31:00
265	2017-070T08:31:00	2017-077T11:24:00
266	2017-077T11:24:00	2017-084T15:15:00
267	2017-084T15:15:00	2017-091T19:06:00
268	2017-091T19:06:00	2017-098T23:00:00
269	2017-098T23:00:00	2017-106T03:03:00

Rev	Start Time (UTC)	End Time (UTC)
270	2017-106T03:03:00	2017-113T03:16:00
271	2017-113T03:16:00	2017-119T14:16:00
272	2017-119T14:16:00	2017-126T01:16:00
273	2017-126T01:16:00	2017-132T12:10:00
274	2017-132T12:10:00	2017-138T23:02:00
275	2017-138T23:02:00	2017-145T10:13:00
276	2017-145T10:13:00	2017-151T21:44:00
277	2017-151T21:44:00	2017-158T09:16:00
278	2017-158T09:16:00	2017-164T20:38:00
279	2017-164T20:38:00	2017-171T07:58:00
280	2017-171T07:58:00	2017-177T19:22:00
281	2017-177T19:22:00	2017-184T06:55:00
282	2017-184T06:55:00	2017-190T18:28:00
283	2017-190T18:28:00	2017-197T05:49:00
284	2017-197T05:49:00	2017-203T17:09:00
285	2017-203T17:09:00	2017-210T04:30:00
286	2017-210T04:30:00	2017-216T15:55:00
287	2017-216T15:55:00	2017-223T03:16:00
288	2017-223T03:16:00	2017-229T14:24:00
289	2017-229T14:24:00	2017-236T01:35:00
290	2017-236T01:35:00	2017-242T12:44:00
291	2017-242T12:44:00	2017-248T23:54:00
292	2017-248T23:54:00	2017-255T10:52:00
293	2017-255T10:52:00	2017-259T09:51:00

Sequence Boundaries:

(Note: Sequences are searchable using the SEQUENCE_ID keyword.)

Sequence	Start Time (UTC)	End Time (UTC)
C27	2001-190T00:00:00	2001-253T00:00:00
C28	2001-252T00:00:00	2001-309T00:00:00
C29	2001-309T00:00:00	2002-014T00:00:00
C30	2002-014T00:00:00	2002-070T00:00:00
C31	2002-070T00:00:00	2002-126T00:00:00
C32	2002-126T00:00:00	2002-189T00:00:00
C33	2002-189T00:00:00	2002-265T22:40:00
C34	2002-265T22:40:00	2002-334T11:36:00
C35	2002-334T11:36:00	2003-039T08:03:00
C36	2003-039T08:03:00	2003-119T04:16:00
C37	2003-119T04:16:00	2003-166T02:17:00
C38	2003-166T01:52:00	2003-222T20:32:00
C39	2003-222T20:32:00	2003-292T16:05:00
C40	2003-292T16:05:00	2004-009T12:25:00
C42	2004-009T12:25:00	2004-051T00:27:00
C41	2004-010T00:00:00	2004-020T00:00:00
C43	2004-051T00:27:00	2004-092T21:28:00
C44	2004-092T21:28:00	2004-135T18:40:00
S01	2004-135T18:40:00	2004-171T21:52:00
S02	2004-171T21:52:00	2004-212T21:32:00
S03	2004-212T21:32:00	2004-256T11:35:00
S04	2004-256T11:35:00	2004-292T09:30:00
S05	2004-292T09:30:00	2004-320T07:49:00
S06	2004-320T07:49:00	2004-351T13:22:00
S07	2004-351T04:22:00	2005-022T10:38:00

Sequence	Start Time (UTC)	End Time (UTC)
S08	2005-022T10:38:00	2005-058T00:36:00
S09	2005-058T00:36:00	2005-099T05:15:00
S10	2005-099T05:15:00	2005-134T02:50:00
S11	2005-134T02:50:00	2005-169T01:34:00
S12	2005-169T01:34:00	2005-212T22:00:00
S13	2005-212T22:00:00	2005-242T21:43:00
S14	2005-242T21:43:00	2005-281T15:57:00
S15	2005-281T15:57:00	2005-316T17:01:00
S16	2005-316T17:01:00	2005-351T14:21:00
S17	2005-351T14:21:00	2006-027T04:03:00
S18	2006-027T04:03:00	2006-070T00:35:00
S19	2006-070T00:35:00	2006-112T05:15:00
S20	2006-112T05:15:00	2006-154T02:39:00
S21	2006-154T02:39:00	2006-198T00:06:00
S22	2006-198T00:06:00	2006-231T22:06:00
S23	2006-231T22:06:00	2006-263T20:22:00
S24	2006-263T20:22:00	2006-295T18:26:00
S25	2006-295T18:26:00	2006-328T16:30:00
B25	2006-295T18:26:00	2006-299T18:11:00
S26	2006-328T16:30:00	2007-005T13:50:00
M27	2007-005T13:50:00	2007-048T10:52:00
S27	2007-005T13:50:00	2007-048T10:52:00
S28	2007-048T10:52:00	2007-087T08:04:00
S29	2007-087T08:04:00	2007-124T22:00:00
S30	2007-124T22:00:00	2007-162T03:10:00

Sequence	Start Time (UTC)	End Time (UTC)
S31	2007-162T03:10:00	2007-195T01:06:00
S32	2007-195T01:06:00	2007-223T23:20:00
S33	2007-223T23:20:00	2007-265T20:51:00
S34	2007-265T20:51:00	2007-304T18:40:00
S35	2007-304T18:40:00	2007-348T16:00:00
S36	2007-348T16:00:00	2008-022T16:35:00
S37	2008-022T13:35:00	2008-047T11:51:00
S38	2008-047T11:51:00	2008-083T01:50:00
S39	2008-083T01:50:00	2008-110T07:18:00
S40	2008-110T07:18:00	2008-152T04:27:00
S41	2008-152T04:27:00	2008-183T19:08:00
S42	2008-183T19:08:00	2008-224T00:20:00
S43	2008-224T00:20:00	2008-257T22:19:00
S44	2008-257T22:19:00	2008-292T20:21:00
S45	2008-292T20:21:00	2008-331T17:55:00
S46	2008-331T17:55:00	2009-009T15:16:00
S47	2009-009T15:16:00	2009-048T12:35:00
S48	2009-048T12:35:00	2009-085T10:05:00
S49	2009-085T10:05:00	2009-125T07:16:00
S50	2009-125T07:16:00	2009-164T04:41:00
S51	2009-164T04:41:00	2009-204T21:51:00
S52	2009-204T21:51:00	2009-237T00:04:00
S53	2009-237T00:04:00	2009-278T04:03:00
S54	2009-278T04:03:00	2009-317T19:21:00
S55	2009-317T19:21:00	2009-356T23:26:00
S56	2009-356T23:26:00	2010-023T15:00:00
S57	2010-023T15:00:00	2010-060T19:04:00
S58	2010-060T19:04:00	2010-095T02:49:00

Sequence	Start Time (UTC)	End Time (UTC)
S59	2010-095T02:49:00	2010-137T13:31:00
S60	2010-137T13:31:00	2010-176T21:10:00
S61	2010-176T21:10:00	2010-211T18:51:00
S62	2010-211T18:51:00	2010-249T06:33:00
S63	2010-249T06:33:00	2010-284T04:17:00
S64	2010-284T04:17:00	2010-328T19:35:00
S65	2010-328T19:35:00	2011-017T08:42:00
S66	2011-017T08:42:00	2011-066T13:02:00
S67	2011-066T13:02:00	2011-115T16:03:00
S68	2011-115T16:03:00	2011-184T11:10:00
S69	2011-184T11:10:00	2011-250T00:48:00
S70	2011-250T00:48:00	2011-320T03:02:00
S71	2011-320T03:02:00	2012-024T22:55:00
S72	2012-024T22:55:00	2012-097T11:47:00
S73	2012-097T11:47:00	2012-170T22:58:00
S74	2012-170T22:58:00	2012-237T18:34:00
S75	2012-237T18:34:00	2012-307T14:30:00
S76	2012-307T14:30:00	2013-013T17:51:00
S77	2013-013T17:51:00	2013-085T13:15:00
S78	2013-085T13:15:00	2013-158T00:45:00
S79	2013-158T00:45:00	2013-226T09:51:00
S80	2013-226T09:51:00	2013-295T23:15:00
S81	2013-295T23:15:00	2013-362T01:47:00
S82	2013-362T01:47:00	2014-072T21:12:00
S83	2014-072T21:12:00	2014-144T10:01:00
S84	2014-144T10:01:00	2014-212T05:09:00
S85	2014-212T05:09:00	2014-279T01:01:00
S86	2014-279T01:01:00	2014-351T03:15:00

Sequence	Start Time (UTC)	End Time (UTC)
S87	2014-351T03:15:00	2015-052T16:52:00
S88	2015-052T16:52:00	2015-121T18:45:00
S89	2015-121T18:45:00	2015-194T13:42:00
S90	2015-194T13:42:00	2015-264T02:48:00
S91	2015-264T02:48:00	2015-329T04:59:00
S92	2015-329T04:59:00	2016-038T00:48:00
S93	2016-038T00:48:00	2016-109T18:43:00
S94	2016-109T18:43:00	2016-178T15:44:00
S95	2016-178T15:44:00	2016-252T10:36:00
S96	2016-252T10:36:00	2016-328T05:43:00
S97	2016-328T05:43:00	2017-028T02:01:00
S98	2017-028T02:01:00	2017-098T21:55:00
S99	2017-098T21:55:00	2017-165T17:15:00

Appendix D: Bibliography

- Alekseeva, G. A., Arkharov, A. A., Galkin, V. D., Hagen-Thorn, E. I., Nikanorova, I. N., Novikov, V. V., Novopashenny, V. B., Pakhomov, V. P., Ruban, E. V., Shchegolev, D. E. 1997. Pulkovo Spectrophotometric Catalog. VizieR Catalog: III/201.
- Bohlin, R.C. and Gilliland, R.L. 2004. Hubble Space Telescope absolute spectrophotometry of Vega from the far-ultraviolet to the infrared. *Astron. J.* 127, 3508–3515.
- Burnashev, V. I. 1985. Catalogue of data on energy distribution in spectra of stars in a uniform spectrophotometric system, Abastumanskaya Astrofiz. Obs. Bull. 59, 83-90. VizieR catalog III/126.
- Glushneva, I. N., Doroshenko, V. T., Fetisova, T. S., Khruzina, T. S., Kolotilov, E. A., Mossakovskaya, L. V., Shenavrin, V. I., Voloshina, I. B., Biryukov, V. V., Shenavrina, L. S. 1998a. Moscow Spectrophotometric Catalog of Stars. VizieR Catalog III/207.
- Glushneva, I. N., Doroshenko, V. T., Fetisova, T. S., Khruzina, T. S., Kolotilov, E. A., Mossakovskaya, L. V., Ovchinnikov, S. L., Voloshina, I. B. 1998b. Sternberg Spectrophotometric Catalog. VizieR Catalog: III/208.
- Gunn, J. E. and Stryker, L. L. 1983. Stellar spectrophotometric atlas, wavelengths from 3130 to 10800 Å. *Astrophysical Journal Supplement Series* 52, 121-153.
- Haemmerle, V. R. & Gerhard, J. H. 2006. Cassini camera contamination anomaly: experiences and lessons learned, *SpaceOps: Earth, Moon, Mars and Beyond*, Rome, Italy, June 19-23, 2006. Available on-line at <http://hdl.handle.net/2014/40797>.
- Hamuy, M., Walker, A. R., Suntzeff, N. B., Gigoux, P., Heathcote, S. R., Phillips, M. M. 1992. Southern spectrophotometric standards. *Astronomical Society of the Pacific*, 104, 533-552.
- Hamuy, M., Suntzeff, N. B., Heathcote, S. R., Walker, A. R., Gigoux, P. 1994. Southern spectrophotometric standards, 2. *Astronomical Society of the Pacific*, 106, 566-589.
- Heck, A., Egret, D., Jaschek, C., Battrick, B. 1984. IUE low dispersion spectra reference atlas. Part 1: Normal stars. IUE low dispersion spectra reference atlas. Part 1: Normal stars, ESA Special Publication: ESA SP-1052 (1984) (ISSN 0379-6566).
- Jamar, C., Macau-Hercot, D., Monfils, A., Thompson, G. I., Houziaux, L., Wilson, R. 1976. UV Bright Star Spectrophotometric Catalog. VizieR On-line Data Catalog: III/39A.
- Karkoschka, E. 1994. Spectrophotometry of the Jovian planets and Titan at 300- to 1000-nm wavelength: The methane spectrum. *Icarus* 111, 174-192.

Karkoschka, E. 1998. Methane, ammonia and temperature measurements of the Jovian planets and Titan from CCD-spectrophotometry. *Icarus* 133, 134-146.

Kharitonov A.V., Tereshchenko V.M., Knyazeva L.N. 1988. Spectrophotometric Catalogue of Stars. *VizieR catalog III/202*, Alma-Ata, Nauka, p. 484

Kurucz, R. L. 1993. Model Atmospheres. *VizieR On-line Data Catalog: VI/39*. Originally published in: 1979 *Ap. J. Supp.*

Ochsenbein, F., Bauer, P., Marcout, J. 2000. The *VizieR* database of astronomical catalogues. *Astron. and Astrophys. Supp.* 143, 23-32.

Porco, C.C., West, R. A., Squyres, S., McEwen, A., Thomas, P., Murray, C.D., DelGenio, A., Ingersoll, A.P., Johnson, T.V., Neukum, G., Veverka, J., Dones, L., Brahic, A., Burns, J.A., Haemmerle, V., Knowles, B., Dawson, D., Roatsch, T., Beurle, K., Owen, W. 2004. Cassini imaging science: instrument characteristics and anticipated scientific investigations at Saturn. *Space Sci. Rev.* 115, 363-497.

Santos, J. F. C., Jr., Alloin, D., Bica, E., Bonatto, C. Spectral Library of Galaxies, Clusters and Stars 2001. *VizieR On-line Data Catalog: III/219*.

West, R. A.; Knowles, B., Birath, E., Charnoz, S., Di Nino, D., Hedman, M., Helfenstein, M., McEwen, A., Perry, J., Porco, C., Salmon, J., Throop, H., Wilson, D. 2010. In-flight calibration of the Cassini imaging science sub-system cameras. *Planetary and Space Science* 58, 1475-1488.