Introduction to the Cassini Imaging Science Subsystem:
Wide Angle Camera

Instrument Overview
====================

The Cassini ISS consists of two fixed focal length telescopes, a
narrow angle camera (NAC) and a wide angle camera (WAC). The WAC is
55 cm long and 35 cm x 33 cm wide, and has a focal length of 200.77
+/- 0.02 mm in the clear filter. The two cameras together have a mass
of 56.9 kg, and sit on the Remote Sensing Palette (RSP), fixed to
the body of the Cassini Orbiter, between the Visual and Infrared
Mapping Spectrometer (VIMS) and the Composite Infrared Spectrometer
(CIRS), and above the Ultraviolet Imaging Spectrometer (UVIS). The
apertures and radiators of both telescopes are parallel to each other.

The WAC has its own set of optics, mechanical mountings, CCD, shutter,
filter wheel assembly, temperature sensors, heaters, and electronics,
the latter of which consists of two parts: the sensor head subassembly
and the main electronics subassembly. The Sensor Head electronics
supports the operation of the CCD detector and the preprocessing of
the pixel data. The Main Electronics provide the power and perform
all other ISS control functions, including generating and maintaining
internal timing which is synchronized to the Command Data System
(CDS) timing of 8 Hz, control of heaters, and the two hardware data
compressors. The Cassini Engineering Flight Computer (EFC) is a
radiation-hardened processor that controls the timing, internal
sequencing, mechanism control, engineering and status data
acquisition, and data packetization.

The optical train of the WAC, a Voyager flight spare, is an f/3.5
refractor with a ~60 microrad/pixel image scale, a 3.5 deg x 3.5 deg
field of view (FOV), and a spectral range from 380 nm - 1050 nm. Its
filter wheel subassembly carries 18 spectral filters: 9 filters on
each of two wheels. This allows for in-line combinations of filters
for greater flexibility. Each wheel is designed to move
independently, in either the forward or reverse direction, at a rate
of 3 positions per second. A homing sensor on each wheel defines a
home wheel position, and wheel positioning can be commanded
absolutely or relatively.

Unlike the NAC, the WAC is not thermally isolated from the RSP. It has
less stringent image quality requirements, so its bulk temperature
control is provided by the pallet.

The temperature of the CCD is controlled by a passive radiator,
directly connected to the focal plane, along with an active
'performance' heater on the CCD to adjust the temperature. The
temperature of the optical elements is controlled by active heaters
positioned along the optical path. These optical elements are kept to
within 1 degree Celsius to maintain camera focus without an active
focusing mechanism. Low expansion invar spacers are also used. The

radiator subassembly also includes two sets of spacecraft-controlled decontamination heaters which are used to minimize deposition of volatile contaminants on either the detector or radiator and to minimize radiation damage to the CCD. All heaters are commandable (ON or OFF) during flight.

Optics

Like the NAC, the focal plane field of view of the WAC is limited by the size of the CCD. However, due to the Voyager optics, the WAC point spread function (PSF) is somewhat larger than a pixel, with a clear filter full width at half maximum (FWHM) of 1.8 pixels. The nominal pixel scale is 59.749 microradians/pixel.

All the optical elements within the WAC are made of either radiation-hardened optical glass (BK7 or lithium fluoride) or fused silica. Antireflection coatings consisting of single layer MgFl2 were deposited on the CCD window and primary optics. A fused silica quartz plug is placed immediately in front of the CCD package to protect the detector against radiation damage and to minimize radiation-induced noise in the images.

The larger field of view of the WAC makes it more susceptible than the NAC to geometric distortions. Measurements of distortion and its dependence on temperature and spectral bandpass in the WAC were made on the ground and in flight. Ground based measurements suggested distortions up to about 3.6 +/- 0.2 pixels in the corners of the CCD, independent of spectral bandpass, in the optics temperature range of -10 degrees C to +25 degrees C. Subsequent observations of the Pleiades and the open cluster M35 showed a consistent distortion parameter of k = -6.27 +/- 0.25, and slight changes in focal length as a function of filter combination. The WAC focal length in the clear filter is 200.77 +/- 0.02 mm. Focal lengths in other filter combinations range from 200.71 mm to 201.22 mm, yielding a range in image radius of 1.27 pixels for a nominal 500 pixel radius object. Thus, individual filter combinations need to be fully calibrated to determine specific focal length. The distortion parameter remains essentially constant in the different filters. In-flight distortion measurements for the WAC are consistent with those taken from the ground: 3.36 pixels in the corners.

Filters

The ISS filter assembly design -- consisting of two filter wheels and a filter changing mechanism -- is inherited from the Hubble Space Telescope WF/PC camera. Each wheel is designed to move independently, in either the forward or reverse direction, at a rate of 2 positions per second in the WAC. A homing sensor on each wheel defines a home wheel position: wheel positioning can be commanded absolutely or relatively.

The WAC filter wheel contains both medium and broad-band filters that

cover the spectral range of the CCD, as well as narrow-band filters for atmospheric studies. The former include the BL1, GRN, RED, IR1, IR2, IR3, and IR4 filters (available on both cameras) as well as VIO and IR5 (WAC only). The latter include the MT2, MT3, CB2 and CB3 filters, used to investigate methane absorption bands and continuum wavelengths. A HAL filter is also included for observing H-alpha emissions from lightning.

The Cassini Imaging Science Team has deliberately duplicated 63% of the filters in both the NAC and WAC. These include seven medium/broadband filters from the blue to the near-IR for spectrophotometry, 2 methane and 2 continuum band filters for atmospheric vertical sounding, 2 clear filters, and the HAL filter.

The clear filter is in the 'home' slot of each filter wheel, since it was deemed that sticking of a filter wheel, should it occur, was most likely to occur in the home position. Typically a clear filter in one wheel is combined with a color filter in the other wheel, though two-filter combinations can also be used. However, with the spare Voyager optics on the WAC, we encountered difficulty in achieving a sharp focus in the near-IR (at which the Voyager vidicon detector was not sensitive). The solution was to place all near-IR filters on one wheel and a special, thin clear filter on the other wheel. As a result of this decision, and because the WAC lacks the UV filters, the only useful 2-filter bandpass in the WAC is IR1-IR2.

Though both cameras are capable of seeing into the near-IR at ~1.0 micron, the wide angle camera is 9 times faster for a given exposure than the NAC and is consequently better equipped to sense this spectral region for either broadband color imaging or atmospheric sounding where the CCD quantum efficiency and solar flux are declining and a large camera throughput is desired, though this benefit is reduced somewhat by the Voyager optical coatings.

Finally, the WAC carries two orthogonal infrared polarizers, IRP0 and IRP90, which can provide intensity and the Stokes parameter, Q, referenced to the principal axes of the polarizers. If the polarizers are oriented parallel or perpendicular to the scattering plane, the information provided by Q is in most cases as informative as that provided by three polarizers because the polarized electric vector is usually aligned parallel or perpendicular to the scattering plane. Estimates of Q referenced to the scattering plane can be made for other orientations but with diminishing precision as the angle between the scattering plane and the polarizer axis approaches 45 degrees at which point the measurement of Q is not useful.

The polarizers are, of course, to be used in combination with other spectral filters and so filter placement was important. In the WAC the 3 broad-band filters at 867 nm, 950 nm, and 977 nm (cutoff filter), and the four narrow-band filters at 727 and 889 nm (methane) and 750 and 940 nm (continuum), are all placed in the same wheel opposite the wheel containing the 2 IR polarizers.

Table 1: ISS WAC Filter Characteristics

```
Filter  Lambda_cen  Lambda_eff  Science Justification
-----------------------------------------------------------------------
VIO     420SP      420         broadband color
BL1     460W       463         broadband color
GRN     567W       568         broadband color
RED     648W       647         broadband color
HAL     656N       656         H-alpha/lightning
MT2     728N       728         methane band, vertical sounding
CB2     752N       752         continuum for MT2
IR1     742W       740         broadband color
IR2     853W       852         broadband color; ring absorption band
MT3     890N       890         methane band, vertical sounding
CB3     939N       939         continuum for MT3,see thru Titan haze
IR3     918W       917         broadband color
IR4     1001LP     1000        broadband color
IR5     1028LP     1027        broadband color
CL1     635        634         wide open, combine w/wheel 2 filters
CL2     635        634         wide open, combine w/wheel 1 filters
IRP0    705        705         IR polarization,see through Titan haze
IRP90   705        705         IR polarization,see through Titan haze
```

Table 2: WAC Two-Filter Bandpasses

```
Filters   lambd_cen   lambda_eff
--------------------------------
IR1-IR2   826         826
```

(All wavelengths in nm. Central wavelengths (lambda_cen) are computed using the full system transmission function. Effective wavelengths (lambda_eff) are computed using the full system transmission function convolved with a solar spectrum. Bandpass types: SP = short wavelength cutoff; W = wide; N = narrow; LP = long wavelength cutoff.)

With the exception of the clear filters and the polarizers, the filters are all interference filters manufactured using an ion-aided deposition (IAD) process which has the effect of making the filters temperature and moisture tolerant, and resistant to delamination. Conventional interference filters have passbands which shift with temperature. The shift can be significant for narrowband filters targeted to methane absorption bands or the H_alpha line. Temperature shifts for IAD filters is typically an order of magnitude or more smaller than for conventional filters and is insignificant over the temperature range (room temperature to 0 degrees C) relevant to calibration and operation of the Cassini cameras.

The WAC infrared polarizers have a 1 mm-thick layer of Polarcor (trademark Corning) cemented between two slabs of BK7-G18 glass. Polarcor is a borosilicate glass impregnated with fine metallic wires. The infrared polarizers have much better performance than the NAC visible polarizers over their range (700 nm - 1100 nm), where the principal transmission is greater than 0.9 and the orthogonal transmission is 0.001 or less.

Shutter

Between the filter wheel assembly and the CCD detector is the shutter
assembly, a two blade, focal plane electromechanical system derived
from that used on Voyager, Galileo and WFPC. To reduce scattered
light, the shutter assembly was put in the optical train `backwards ,
with the unreflective side towards the focal plane. Each blade moves
independently, actuated by its own permanent magnet rotary solenoid,
in the sample direction: i.e., keeping the blade edge parallel to the
columns of the CCD. The shutter assembly is operated in 3-phases: open
(one blade sweeps across the CCD), close (the other blade sweeps
across the CCD to join the first), and reset (both blades
simultaneously sweep across the CCD in the reverse direction to the
start position).

There are 64 commandable exposure settings which can be updated during
flight if so desired. These correspond to 63 different exposure times,
ranging from 5 milliseconds to 20 minutes, and one `No Operation
setting. The shortest nonzero exposure is 5 msec. In the ISS flight
software, the time tag on the image is the time of the close of the
shutter. Because of mechanical imperfections in the shutter mechanism,
there is a difference between the commanded exposure time and the
actual exposure time, and a gradient in exposure time across the CCD
columns. At an operating temperature of 0 degrees C, the mean
differences in the WAC for commanded exposure times of 5, 25 and 100
ms were measured to be 0.15, 0.39 and 0.07 ms, respectively. In all
cases the actual exposure times are less than the commanded times.
There is also a small temperature dependence to these shutter offsets.

The 1024th column is illuminated first in both cameras. In the WAC,
this column is illuminated for ~ 0.1 msec longer than the first
column. This value is independent of exposure time and reasonably
independent of temperature. The expected precision or repeatability
of an exposure (equal to the standard deviation of actual exposure
durations measured at any one location on the CCD in ground tests) is
</= 0.04 msec for the WAC. Corrections for the mean and the
spatially-dependent shutter offsets are incorporated into the Cassini
ISS calibration software (CISSCAL). The shutters were tested for
light leak, of which a small signal was detected in the WAC. It
produced 1 DN (12 electrons) or less at a level of 10,000 times full
well incident on the closed shutter.

Detector

The CCD detector used in the Cassini ISS was manufactured by Loral,
packaged by JPL, and employs three phase, front-side-illuminated
architecture. The imaging area -- the region on which light is focused
-- is a square array of 1024 x 1024 pixels, each 12 microns on a side.
The CCDs on both cameras were packaged, hermetically sealed and
fronted by a fused silica window.

The CCD's response to light is determined by the spectral dependence

of each pixel's quantum efficiency: i.e., the number of electrons released in the silicon layer for each photon incident on it. In front-side-illuminated CCDs (like that in the Cassini ISS), the overlying polysilicon gate structures don t transmit UV light. To achieve the required UV response, a UV-sensitive organic fluorescent material called lumogen was vacuum-deposited onto the CCD at 80 degrees C after it was bonded. In this 0.6 micron layer, UV photons are converted into visible photons in the 540 to 580 nm range that readily penetrate the silicon below. Under vacuum conditions, the lumogen layer would tend to evaporate when CCD temperatures reached 60 degrees C. For this reason, the CCD sealed packages were back-filled with inert argon gas to a half atmosphere pressure. All flight candidate CCDs were coated with lumogen before the two flight CCDs were chosen and assigned to each camera. Hence, despite the fact that the WAC optics don t transmit in the UV, the WAC CCD is also coated with lumogen.

The efficiency of a CCD in the near-IR depends on its thickness, or more precisely on the thickness of the very thin, high purity silicon layer which is epitaxially grown over a thicker (~ 500 micron) substrate. It is the photons absorbed in the epitaxial layer that are converted into the signal electrons that are subsequently collected and sampled. Nearly all of the near-IR photons actually penetrate beyond the epi layer and create charge in the substrate. However, the purity contrast between the substrate and the epi layer prevents substrate-generated charge from entering the epi layer and being collected. Thus, the 1100 nm quantum efficiency is essentially the fraction of incident flux which is absorbed in the thin layer of pure silicon: the thicker the epi layer, the higher the infrared sensitivity. However, the thicker this layer, the lower the spatial resolution. A compromise was made in the manufacture of the CCD to yield some response near 1100 nm while maintaining high spatial resolution. The epi layer is 10 - 12 microns thick on Cassini and results in a quantum efficiency (QE) of ~1% at 1000 nm.

A compromise involving the near-IR response was also made in choosing the CCD operating temperature. At Saturn, this temperature is -90 +/- 0.2 degrees C and is a compromise between yielding an acceptably low dark current (</= 0.3 e-/sec/pixel) and maintaining a reasonable near-IR response (which is diminished at low temperatures). CCD thermal control is achieved by means of balance between passive radiation to space, which alone would maintain the CCD below its operating temperature at Saturn, and active heater control. The radiator of each camera also supports a decontamination heater (35 watts in all) that can heat the CCD to +35 degrees C to reduce the deposition of volatile contaminants on either the detector or the radiator. (Because damage to the CCD due to cosmic rays can be annealed at elevated temperatures, the CCD operating temperature during cruise, when data were not being collected, was maintained at 0 degrees C to minimize such damage.)

The detector system includes an unilluminated region 8 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 up into the serial register and out to the signal chain, the read-out continues for 8 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 [PORCOETAL2004].)

Scientific Objectives
======================

See [PORCOETAL2004] for an in-depth description of Cassini ISS science objectives.

Camera Operation
=================

Operational States

The ISS has three operational power states: On, Sleep and Off. In the On state, the cameras are Active or Idle. In this state, both the spacecraft replacement heaters and ISS decontamination heaters are off. The camera software has active control over the performance heaters to set appropriate operating temperatures for the optics and CCD detector. The Active sub-state is entered to collect science data as well as for calibration and maintenance activities. Command execution in the active state includes science data readout, filter wheel movement, shutter movement, activation of light flood and calibration lamps, and other high power consuming activities. Idle is a background state in which no commands are executing. When the camera is in Idle, uploads can be processed, real-time and 'trigger' commands can be accepted from the CDS, and macros can be stored. The execution of any command sends the camera into the Active state. The camera always returns to Idle state after completing a command sequence. In the WAC, peak power consumption during active imaging is 19.4 watts.

The ISS Sleep state is a non-data taking low power state that is used when no activity is planned for an extended period of time. During this state, the sensor head and main ISS electronics are drawing power, and the optics and CCD heaters are on to maintain operating temperature limits. Spacecraft controlled replacement heaters are off. The decontamination heaters may be used, if necessary. In Sleep, the WAC consumes 16.4 watts.

In OFF, no power is drawn by the ISS. The spacecraft controlled replacement heaters and ISS decontamination heaters may be turned on when necessary. The replacement heaters keep the ISS within allowable flight nonoperating temperature limits and the decontamination heaters

can be used to provide for CCD protection from the radiation environment and from the condensation of volatiles. In this state, the WAC consumes 4.5 watts.

Detector Modes

The CCD has the capability of being commanded to operate in full mode (i.e., 1x1) or either 2x2 or 4x4 on-chip pixel summation modes. The latter two modes are used for either enhancing signal-to-noise and/or decreasing the data volume and/or read-out time at the expense of spatial resolution. The full well of the CCD is roughly 120,000 e-/pixel. Four gain states are available: for imaging faint objects (high gain, Gain 3) and bright objects (normal gain, Gain 2), and to match the output of the 2x2 (Gain 1) and 4x4 (Gain 0) full wells. The summation well can hold only 1.6×10^6 electrons; this corresponds to full well with 4X4 summing. However, the relation between number of electrons in the signal and the digital data numbers (DN) into which the signal is encoded starts to become nonlinear above 10^6 electrons because at this signal level, the on-chip amplifier becomes non-linear. For this reason, in the lowest gain state (Gain 0), the full scale signal is set to correspond to $\sim 10^6$ electrons at 4095 DN.

Table 3: WAC Gain States

Gain State	e-/DN	Notes
0	211 +/- 16	Designed for 4x4 summation mode
1	85 +/- 7	Designed for 2x2 summation mode
2	28 +/- 1	Normal gain; used in 1x1 summation mode
3	12 +/- 1	Used in 1x1; chosen to match read noise

The capability also exists within the ISS to reduce the effect of blooming, the phenomenon whereby a highly overexposed pixel can spill electrons along an entire column of pixels, and sometimes along a row, when the full well of the CCD is exceeded. The default camera setup has anti-blooming on, with the option to turn it off. Anti-blooming mode is achieved by applying an AC voltage to the chip, forcing excess electrons into the silicon substrate. An undesirable side effect of this action is to pump electrons into traps in the silicon at the expense of electrons in adjacent pixels. For long exposures this produces bright/dark pixel pairs.

Camera Commanding

The acquisition of images can be accomplished in several ways. Individual NAC or frames may be acquired, or the NAC and WAC can be used in simultaneous mode, called BOTSIM (for 'both simultaneous'). The entire event, which is called a framing event and requires a total duration called a 'framing time', is broken down into two steps: the prepare cycle and the readout cycle.

The prepare cycle is used to alter the state of the ISS, step the

filter wheels, perform heater operations, light flooding, and other functions required to prepare for an exposure. It also includes the exposure time. The prepare cycle is constructed from a series of quantized windows of time in which specific functions are assigned to occur.

During the prepare cycle, the shutter blades are reset from the previous exposure and the filter wheels are moved into position. Because simultaneous motion of each filter wheel requires more power than the ISS was allocated for peak operation, all filter wheels NAC and WAC -- are moved separately. Windows of quantized duration are set aside for the motion of each filter wheel. Next, the CCD is prepared for exposure to light. This preparation begins with a wait; the duration of the wait is chosen to ensure that the shutter will close exactly at the end of the prepare cycle. After the wait, a light-flood fills the wells of the CCD to many (~ 50) times saturation, followed immediate by a read out. The entire light-flood/erase event takes 950 msec and has the effect of erasing any residual image of previous exposures from the CCD. Within 5 msec of the end of the light-flood/erase event, the shutter is opened for the commanded duration. (For dark frames, this duration is set to zero.) The image is tagged with the time of shutter close.

During a BOTSIM, the prepare cycle is lengthened to include time to prepare both NAC and WAC. The NAC is prepared first; then the WAC is prepared so as to avoid simultaneous movement of any of the 4 filter wheels. If the NAC and WAC exposure times are different, the exposures begin in a staggered fashion so that the NAC and WAC shutters are closed simultaneously. There are 63 discrete commandable exposure times which are accommodated within 13 discrete prepare cycle windows.

During the following readout window, the CCD is read out, the data are encoded and/or compressed, and the results are packetized. For any of the 6 individual CDS pickup rates, there are 4 discrete readout windows for each camera. The readout window is scaled by the CDS pickup rate giving 24 actual readout windows per camera and 96 actual BOTSIM readout windows.

Prepare times and readout times are chosen before uplink. The prepare cycle is completely determinate; the readout time required to fully read out an image is not. The required readout time during the image event will depend on the amount of data being read out of the CCD, and the CDS pickup rate or on the line readout rate from the CCD, whichever is slowest. If the data volume in the image was underestimated and the required readout time exceeds the commanded readout time, the camera will cease reading out part way through an image and lines will be lost. For this reason, a great deal of effort has gone into the amount of data returned for different scenes and choices of compression parameters.

The ISS can collect pixel (image) data, engineering data and status data, and packetize them with appropriate header information as either science telemetry packets (which include all types of data) or housekeeping packets (which only include engineering and status data).

The latter are returned alone when ISS is in an ON power state but not actively taking images. The frequency with which housekeeping packets are collected is 1 packet/sec and is programmable in flight. The amount of housekeeping data that gets sent to the ground is determined by the rate at which CDS picks up such packets and is currently 1 housekeeping packet every 64 seconds.

Data paths

The analog to digital (A/D) conversion happens right as the analog signal is read out from the chip, after it has passed through the on-chip amplifier. Data from the ADC are encoded to 12-bit data numbers (DN), giving a dynamic range of 4096. However, they are stored as 16-bits: the upper 4-bits are all 1 s. The ISS flight software masks the upper 4 bits when doing calculations. Compression and conversion functions are performed after the electrons are converted to DN. The next juncture is a choice of data conversion (from 12 to 8 bits) or no data conversion. Unconverted data can then proceed to a lossless compressor or undergo no compression at all. Converted data can undergo no compression or lossless or lossy compression. From there, the data are placed on the Bus Interface Unit (BIU), where they are ultimately picked up by the Command Data System (CDS) and sent to the Solid State Recorder (SSR) where they are stored as 16-bit data.

Data Compression

Serious constraints are imposed on imaging of the Saturn system by the limited storage volume on the spacecraft's SSR, and by the limited communication bandwidth back to Earth. In order to make the most effective use of these resources, the Cassini imaging system includes the capability to convert the data from 12 bits to 8 bits (called data conversion), and also to perform either 'lossless' or 'lossy' image compression. Data conversion, and both lossless and lossy compression, are implemented in hardware.

Conversion to 8 bits

Two sub-options are available for 8-bit conversion. One is a variant on conventional 'square root' encoding. In such encoding, a look-up table (LUT) is used to convert the original data values to 8-bit values. The output 8-bit values are related to the input values in a non-linear fashion, typically scaling with the square root of the 12-bit value. This non-linear scaling more closely matches the quantization level to the photon shot noise so that the information content is spread more evenly among the 256 levels. (The Cassini 12-to-8 bit conversion table is provided with the calibration data volume.) It differs somewhat from pure square-root encoding, having been designed for the known noise properties of the Cassini cameras to distribute quantization-induced errors uniformly across the dynamic range of the system. The look-up table is stored in ROM

within the cameras' memory and cannot be altered in flight; choice of ON or OFF is commandable in flight.

The other sub-option is conversion from 12 bits to the least-significant 8 bits LS8B). This type of conversion is useful for reducing the data volume of images taken of very faint targets, such as diffuse rings or the dark side of Iapetus, which generally do not yield large signal levels and can be encoded to the lowest 8 bits.

Lossless Compression

Both converted (8-bit) and unconverted (12-bit) data can be lossless compressed. The ISS lossless hardware compressor is based on Huffman encoding, a high efficiency, numerical encoding scheme in which the length of the bit sequence used to encode a given number is chosen based on the frequency of occurrence of that number. In ISS lossless compression, each compressed image can be reconstructed on the ground with no loss to the information content of the image, provided the image entropy does not exceed the threshold where 2:1 compression is achieved. Scenes with low entropy will have compression ratios higher than 2:1; scenes with high entropy will never compress greater than 2:1, but the ends of lines will be truncated so that the total amount of data returned in a pair of lines of the image never exceeds the total number of bits for a single uncompressed line. The truncation scheme has been designed so that the truncation alternates -- i.e., every other line -- from one line to the next, on the right (large sample number) side of the image. If the data loss is great, it can in principle result in the complete loss of every other line. In either case, with this scheme, information (though reduced in spatial resolution) can be retained across the image.

Lossy Compression

Imaging sequences requiring larger compression ratios than can be achieved with the lossless compressor may instead be more strongly compressed using the camera's lossy compression circuitry. This capability requires that the data have been converted to the 8-bit form. Consequently, data conversion must be employed first before the data are sent to the lossy compressor. Compression is implemented by a pair of specialized signal processing chips which perform a variation on the familiar JPEG (Joint Photographic Experts Group) compression algorithm used in many image transfer and storage applications. The JPEG algorithm operates by selectively removing information from an image, particularly at high spatial frequencies. Lossy-compressed images thus tend to have reduced detail on fine scales.

For More Information
=====================

More information regarding the camera design, operation, imaging and compression modes, and image calibration can be found in

[PORCOETAL2004]. Additional discussion of calibration can also be found in the documentation for the calibration volume of this data set. "

```
  END_OBJECT                = INSTRUMENT_INFORMATION

  OBJECT                    = INSTRUMENT_REFERENCE_INFO
    REFERENCE_KEY_ID        = "PORCOETAL2004"
  END_OBJECT                = INSTRUMENT_REFERENCE_INFO
END_OBJECT                  = INSTRUMENT
END
```