

Orbiting Carbon Observatory – 2 & 3 (OCO-2 & OCO-3)



Level 1B Algorithm Theoretical Basis

Version 3.0, rev 0, for OCO-2 v11 & 11R data and OCO-3 v10/10.4 datasets

February 6, 2024

National Aeronautics and
Space Administration



Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

ORBITING CARBON OBSERVATORY (OCO) – 2 & 3

LEVEL 1B

Theoretical Basis Document

Robert Rosenberg	Jet Propulsion Laboratory
Lars Chapsky	Jet Propulsion Laboratory
Graziela R. Keller Rodrigues	Jet Propulsion Laboratory
Richard Lee	Jet Propulsion Laboratory
Yuliya Marchetti	Jet Propulsion Laboratory
Aronne Merrelli	University of Michigan
Fabiano Oyafuso	Jet Propulsion Laboratory
Gary Spiers	Jet Propulsion Laboratory
Shanshan Yu	Jet Propulsion Laboratory
Jia Zong	Jet Propulsion Laboratory
Abhishek Chatterjee	Jet Propulsion Laboratory
Vivienne Payne	Jet Propulsion Laboratory
Greg Osterman	Jet Propulsion Laboratory
David Crisp	Jet Propulsion Laboratory (Retired)
Annmarie Eldering	National Institute of Standards and Technology
Randy Pollock	Jet Propulsion Laboratory (Retired)

Document History

Version	Revision	Date	Description/Comments
1.0	0	April 14, 2014	Initial version of OCO-2 L1B ATBD
1.1	0	December 30, 2014	Updated in preparation for L1B public release
1.2	0	March 17, 2015	Version submitted to L1B review panel
1.2	1	March 30, 2015	Version delivered to DAAC alongside L2 release
1.3	0	June 10, 2015	Version updated for release with version 7 and 7r data
1.4	1	October 14, 2017	Version updated for release with version 8 and 8r data
1.3	1	January 2, 2019	Added OCO-3-specific material
2.0	0	May 15, 2021	Version updated for release of Version 10 and 10r OCO-2 data and OCO-3 vEarly and vEarlyR data
3.0	0	February 6, 2024	Version updated for release of Version 11 and 11r OCO-2 data and OCO-3 v10 data

The research described in this document was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Copyright 2024 California Institute of Technology. U.S. Government sponsorship acknowledged.

TABLE OF CONTENTS

1	Scope of This Document and Updates.....	1
1.1	Updates since Version 1.3, Rev 1 (2 January 2019).....	1
2	Overview of the OCO-2 and OCO-3 Measurement Approach.....	3
2.1	Instrument Characteristics	4
2.1.1	Spectral and Geometric Characteristics.....	4
2.1.2	Polarization Characteristics	7
2.2	The OCO-2 and OCO-3 Onboard Calibration Systems.....	7
2.3	Types of On-Orbit Calibration Measurements	10
2.4	Data Product Delivery.....	13
3	Radiometric Calibration.....	15
3.1	Bad Pixels and Bad Samples.....	16
3.2	Bad Pixel Map.....	17
3.2.1	Bad Pixel Map Usage	18
3.2.2	Summed Mode Pixel Map Processing.....	18
3.2.3	Bad Pixel Map/Footprint Relationship.....	18
3.2.4	Bad Pixel Mitigation Calculations.....	19
3.3	Dark Subtraction	19
3.4	Radiometric Gain	20
3.5	Radiometric Gain Corrections	21
3.5.1	Fast Degradation.....	21
3.5.2	Slow Degradation	21
3.5.3	OCO-2 Gain Degradation and Zero-Level Offset Correction Methodology ..	23
3.5.4	OCO-3 Gain Degradation.....	26
3.6	Zero-level Offset.....	27
3.7	Signal-to-Noise Ratio (SNR)	29
3.7.1	Calculation of SNR.....	29
3.7.2	Development of SNR Coefficients.....	30
3.7.3	In-Flight Updates of SNR Coefficients	30
3.8	Bad Sample List.....	31
3.9	Radiometric Footprint Bias and “Streak Flats”	31
4	Spectral Calibration	36
4.1	Spectral Dispersion Coefficients.....	36
4.2	Instrument Line Shape	37

5	Geometric Calibration.....	39
5.1	Pointing and Geolocation.....	39
5.2	Radiance Discontinuities Due to FPA Rotation—The Clocking Correction.....	40
5.2.1	Clocking Algorithm Description	46
5.2.2	Validation of the Clocking Correction	51
6	Polarimetric Calibration.....	58
6.1	Treatment of Polarization	58
6.2	OCO-2 Polarization	59
6.3	OCO-3 Pointing Mirror and Polarization	60
6.3.1	PMA Movement in Non-nadir Modes.....	60
6.3.2	PMA Polarization Angle	62
6.3.3	PMA Polarization Angle and ISS β Angle.....	63
7	Cosmic Ray Artifacts.....	66
8	Ancillary Radiometric Product (ARP) Files.....	68
9	Other Resources	69
10	References.....	70

LIST OF FIGURES

Figure 2-1.	The OCO-2/OCO-3 instrument showing the major optical components and optical path.....	5
Figure 2-2.	(a) The illumination and readout approach used for the OCO-2 and OCO-3 FPAs.	6
Figure 2-3.	(Left) OCO-2’s Onboard Calibrator (OBC) is integrated into the telescope baffle and can be rotated to place diffusers in front of the aperture. (Right) A drawing of the OBC showing the aperture cover, solar diffuser, and the calibration paddle in the open, closed, and solar diffuser postions.....	8
Figure 2-4.	A schematic diagram illustrating the process for generating Level 1B files from Level 0 files.....	13
Figure 3-1.	Overview of the ground processing data flow for the radiometric calibration process used to generate L1B data.....	15
Figure 3-2.	Example of a term in the OCO-2 dark subtraction (dn_{raw}).....	20
Figure 3-3.	OCO-2 relative solar cal radiometric response for the ABO2 (blue), WCO2 (green), and SCO2 (red) channels as a function of time.	22
Figure 3-4.	Slow loss of solar calibration signal in the (a) ABO2 channel, (b) WCO2 channel, and (c) SCO2 channel.	22

Figure 3-5. Lunar calibration time series after correction for observing geometry, including polarization effects, and icing in the ABO2 (top), WCO2 (middle) and SCO2 (bottom) channels indicate that the irreversible loss of radiometric throughput of the instrument is ~0.5%/year in the ABO2 channel (blue) and < 0.1%/year in the WCO2 (green) and SCO2 (red) channels.....	23
Figure 3-6. The zero-level offset, as measured in unilluminated rows in the shadow of the OCO-2 slit mask, early (left) and later (right) in a decon cycle.	24
Figure 3-7. A-Band gain degradation factors from early (top) and late (bottom) in a fast degradation cycle.	25
Figure 3-8. Time series of gain degradation factors for a single sample in each band.....	25
Figure 3-9. Differences between radiances from OCO-2 v8 and v7.	26
Figure 3-10. Trends of radiance relative to prelaunch for the OCO-3 ABO2, WCO2, and SCO2 channels with different colors indicating different lamps.....	27
Figure 3-11. OCO-2 A-band ZLO correction factors near the beginning of (cyan) and deep into (blue) a fast degradation cycle.	28
Figure 3-12. Example of OCO-2 signal-to-noise coefficients.	30
Figure 3-13. The two panels show the orientation of the slit with respect to the track direction at two different moments during an orbit.	32
Figure 3-14. Footprints 1, 3, 5, and 7 and 2, 4, 6, and 8 from seven frames around the streak flat frame nearly overlap.	33
Figure 3-15. Selected continuum regions: ABO2: 0.7707–0.7711, WCO2: 1.60567–1.60593, SCO2: 2.06015–2.0603 microns.....	34
Figure 3-16. Factors (colored circles) derived from the streak-flat analyses for each footprint, of each band, for three time periods, used to divide pre-launch gains.....	35
Figure 4-1. An example of the conversion of the focal plane column to wavelength scale.	37
Figure 4-2. Example OCO-2 ILS profiles for the O ₂ A band at three different spectral pixel indexes.	38
Figure 5-1. (A) topographic elevation variations near the Lauder TCCON station. (B) Estimates of V8 XCO ₂ retrieved from OCO-2 observations collected near the Lauder station and in other regions with significant topography showed anomalous variations that were strongly correlated with topographic slopes. (C) Correcting the pointing error in the V9 product reduces XCO ₂ biases by more than half, illustrating the impact on XCO ₂ products of continual improvements to L1b products.	40

Figure 5-2. (a) Images of atmospheric spectral from the OCO-2 ABO2 (top), WCO2 (middle) and SCO2 (bottom) channels showing the rotation, or “clocking” of the FPA columns with respect to the dark O ₂ and CO ₂ absorption lines. The 8 spatially-summed footprints in each channel are shown from bottom (footprint 1) to top (footprint 8) in blue and red shading. The clocking is most obvious is in the ABO2 channel, where the O ₂ lines are tilted slightly counterclockwise and in SCO2 channels, where the CO ₂ lines are tilted clockwise. The columns where the clocking corrections are applied in each channel are highlighted with red circles. (b) Same as panel (a) for the OCO-3 instrument. For OCO-3, only the ABO2 channel needs a clocking correction because the WCO2 and SCO2 channels are well aligned.....	41
Figure 5-3. The onboard software "clocking" scheme used to sample the same spatial information (vertical dimension) across the entire spectral range (horizontal) for a rotated FPA.	42
Figure 5-4. Schematic illustrating the clocking for the OCO-2 O ₂ A band along with the current color slice selection.	43
Figure 5-5. The color slices used in the clocking correction for the O ₂ A band (shown in Figure 5-4) are aggregated into five “groups.”.....	44
Figure 5-6. Schematic illustrating the clocking for the WCO2 band along with the current color slice selection.	44
Figure 5-7. The color slices used in the clocking correction for the WCO2 band (shown in Figure 5-6) are aggregated into four “groups.”.....	45
Figure 5-8. Schematic illustrating the clocking for the SCO2 band along with the current color slice selection.	45
Figure 5-9. The color slices used in the clocking algorithm for the SCO2 band (shown in Figure 5-8) are aggregated into two “groups.”.....	46
Figure 5-10. O ₂ A band radiances measured for a few select color slices as a function of row index, which serves as a proxy for actual spatial variation within a scene.....	48
Figure 5-11. Example of a clocking correction function for the O ₂ A band.....	49
Figure 5-12. Weighting functions for each group of color slices for the OCO-2 O ₂ A band.	50
Figure 5-13. An example OCO-2 O ₂ A spectrum, both before and after clocking correction.	51
Figure 5-14. Orbit used in discussion of the clocking correction	52
Figure 5-15. Cumulative distribution function of magnitude of the correction for orbit 1600.....	52
Figure 5-16. Correction correlation matrix for the O ₂ A band.....	53
Figure 5-17. Correction correlation matrix for the WCO2 band.	54
Figure 5-18. Correction correlation matrix for the SCO2 band.	55
Figure 5-19. Right singular vectors corresponding to eight largest singular values from violet (largest) to red (smallest) for the O ₂ A-band.	56
Figure 5-20. First eight right singular vectors for the WCO2 band.	56
Figure 5-21. First twelve right singular vectors for the WCO2 band.	57

Figure 6-1. Geometry of the reference plane for polarization, i.e., the local meridian plane.	59
Figure 6-2. Polar coordinate view of the PMA actuator angles, with the PMA elevation plotted as the radial coordinate, for positive elevation (left) and negative elevation (right).	62
Figure 6-3. Top: View of the polarization angle (using the approximation in the equation above), folded to the range 0–90 degrees, using the dual polar coordinate axis (same as Figure 6-2).	63
Figure 6-4. Example paths in the PMA actuator angle space for glint observations, at a range of ISS β angles.	64
Figure 6-5. SNR calculation for Cox-Munk surface model as a function of Solar Zenith angle and polarization angle (background filled contours), with over plotted points for various paths followed by the PMA for glint mode observations.	65
Figure 7-1. The number of spectral samples in the OCO-2 O ₂ A band that are not contaminated by 10-sigma cosmic ray events clearly shows the geographical extent of the South Atlantic Anomaly (SAA).	66
Figure 7-2. An example showing the impact of a radiation event on spectra.	67

LIST OF TABLES

Table 2-1. Summary of OCO-2 instrument performance.	4
Table 2-2. Summary of OCO-3 instrument performance.	4
Table 2-3. Summary of inflight calibration measurements.	10
Table 3-1. History of OCO-2 bad pixel map uploads.	17
Table 3-2. History of OCO-3 bad pixel map uploads.	17
Table 3-3. OCO-2 footprint limits at the long wavelength end.	18
Table 3-4. Bad pixel mitigation.	19
Table 3-5. OCO-2 and OCO-3 <i>MaxMS</i> values.	29
Table 3-6. Minimum measurable signal (<i>MinMS</i>) values.	30
Table 5-1. Glossary of symbols.	47

1 Scope of This Document and Updates

This algorithm theoretical basis document (ATBD) describes the methods used to convert measurements returned by the Orbiting Carbon Observatory-2 (OCO-2) and Orbiting Carbon Observatory-3 (OCO-3) spectrometers from time-ordered Level 1A products, expressed as data numbers (dn), into Level 1B (L1B) data products, consisting of radiometrically, spectrally, and geometrically calibrated radiance spectra. An overview of the measurement approach and OCO-2 and OCO-3 instruments is given in Section 2. The reader is encouraged to consult the published articles referenced in this document for details concerning the analyses of the instrument pre-launch and on-orbit performance for OCO-2 and OCO-3.

The OCO-2 and OCO-3 Data User's Guide discusses OCO-2 mission objectives and measurements. Please refer to that document for details of the data collection and naming convention. The algorithms used to derive L2 products are described in their respective ATBDs.

1.1 Updates since Version 1.3, Rev 1 (2 January 2019)

The algorithm for producing OCO-2 L1B data has undergone a series of updates since the first official release, the Version 6 product, which was released in December 2014. The first major reprocessing campaign for OCO-2 was conducted in 2015, yielding the Version 7 product. That product included improved gain corrections for spectral samples that had developed bad pixels. It also included a preliminary correction for gain degradation in the ABO2 channel, but all of the observed loss in solar cal signal was attributed to the instrument. This approach led to radiometric calibration errors that increased to about 6% in this channel by mid-2017. The v7 L1B product also included no gain degradation corrections for the WCO2 or SCO2 channels.

The second major reprocessing effort for OCO-2 produced the Version 8 (v8R) product, which was delivered in 2017. In this delivery, solar observations acquired after each instrument decontamination campaign (decon) were used to trend the long-term degradation of the solar diffuser and other parts of the optical system. The v8 product included a simple, spectrally- and spatially- constant correction to the throughput of the WCO2 and SCO2 channels, based on observed changes in the median solar calibration results. All of the long-term throughput changes were attributed to reductions in the throughput of the solar diffuser. In addition to these changes, the v8 product included a correction for a small (0.2–0.9 % of the continuum level) spectrally dependent “zero level offset” in the measured ABO2 radiances that is strongly correlated with ice accumulation on the FPA. A small change in the spectral dispersion coefficients was also adopted in v8.

In October 2018, the geometric calibration was updated to correct a small error in the boresight pointing and geolocation. This change resulted in the production of the Version 9 (v9) product. This was the only significant change in v9.

Version 10 was the first version of the L1B algorithm used to support both OCO-2 and OCO-3. For OCO-2, the v10 product included an improved treatment of radiometric degradation that used lunar calibration observations to more clearly separate throughput degradation in the solar diffuser from that in the optical paths of the spectrometers used for science measurements. Observations through 2019 estimated that the throughput of the ABO2 channel was decreasing by 0.61 %/year, while changes in the throughput of the two CO₂ channels could be resolved. The

v10 product also implemented another small refinement in the spectral dispersion coefficients and the identification of additional outlier spectral samples.

For OCO-3, the v10 algorithm was used to generate the “vEarly” product, which was released in July 2020. In this version, the absolute radiometry is poorly constrained because it is based entirely on calibration observations using the primary onboard lamp, whose output has changed substantially since in-orbit checkout. The next version of calibration (v10.3R) uses the secondary and tertiary onboard lamps, which are used less often and have changed much less. Efforts are also ongoing to better constrain the radiometry using observations of the Moon and surface calibration targets.

This version of the ATBD supports the Version 11 (v11) release. A longer record of lunar observations showed ABO2 gain degradation to be slowing, so instead of a linear decay an exponential function was implemented. Reduced uncertainty and removal of a spurious seasonal trend also allowed very subtle linear decays to be implemented in the WCO2 and SCO2 channels. In addition, the ABO2 instrument line shape was refined by accounting for stray light in the preflight laser and heliostat measurements. The “wings” were reduced while the “core” was largely unchanged. Another minor adjustment was to derive the background term of the noise model from inflight dark measurements instead of preflight as in previous versions. Finally, as with previous versions, additional outlier spectral samples were identified and spectral dispersion offset and stretch were updated.

2 Overview of the OCO-2 and OCO-3 Measurement Approach

The OCO-2 and OCO-3 missions carry nearly identical, three-channel, high-resolution imaging grating spectrometers. These instruments make coincident, co-boresighted measurements of reflected sunlight in the short-wave-infrared (SWIR) CO₂ bands centered near 1.61 and 2.06 μm and in the near-infrared (NIR) molecular oxygen (O₂) A-band at 0.765 μm (Crisp et al. 2017). Simultaneous measurements from these three spectral regions are combined to define a single “sounding.” Each sounding is analyzed with remote-sensing retrieval algorithms to produce an estimate of XCO₂ for the atmospheric path between the Sun, the reflecting surface, and the instrument aperture. Measurements of absorption within the weak CO₂ band near 1.61 μm and the stronger CO₂ band near 2.06 μm provide information about the CO₂ column abundance. Measurements of absorption in the O₂ A-band provide constraints on the surface pressure and uncertainties in the atmospheric optical path length introduced by cloud and aerosol scattering and pointing errors. The OCO-2 and OCO-3 spectrometers collect spectra in eight contiguous spatial footprints across a narrow field of view (FOV) at 1/3 second intervals, yielding 24 soundings per second. OCO-2 footprint areas are < 3 km² along a < 10 km wide ground track, and OCO-3 footprint areas are < 4 km² along a < 14 km wide ground track.

In addition to XCO₂, the OCO-2 and OCO-3 spectrometers return precise estimates of solar-induced chlorophyll fluorescence (SIF). This fluorescence must be quantified and corrected in the O₂ A-band channel to ensure accurate estimates of XCO₂ (Frankenberg et al. 2012, 2011a). SIF is also a functional proxy for terrestrial gross primary productivity (GPP; Frankenberg et al. 2011b; Guanter et al. 2012; Joiner et al. 2013; Köhler et al. 2015; Sun et al. 2017, 2018; Magney et al. 2019). SIF-based estimates of GPP are being combined with XCO₂ in atmospheric inverse models to provide new constraints on CO₂ uptake by the land biosphere and new insights into carbon-climate feedbacks (Liu et al. 2017; Palmer et al. 2019; Crowell et al. 2019).

The OCO-2 spacecraft flies in formation with the 705-km Afternoon Constellation (A-Train). This 705-km altitude, 98.8-minute, Sun-synchronous orbit has a 98.2-degree inclination, a 1:36:30 PM mean ascending equator crossing time, and a 16-day (233 orbit) ground track repeat cycle. OCO-2 follows a ground track that crosses the equator 217.3 km east of the World Reference System-2 ground track that was followed by the NASA Aqua platform. This orbit was originally chosen such that OCO-2 overflies the CloudSat radar and CALIPSO lidar ground footprints. Those two spacecraft left the A-Train to form the C-Train in September of 2018. OCO-2 remained in its A-Train orbit because its early afternoon mean local time is better suited for acquiring observations of the absorption of reflected sunlight by CO₂ and O₂ because the Sun is high, maximizing the available signal. It also facilitates coordinated calibration and validation campaigns with other A-Train instruments, and synergistic use of OCO-2 data with that from other A-Train platforms. For routine science operations, OCO-2 points the spectrometer’s FOV at the local nadir or near the “glint spot” where sunlight is specularly reflected from the surface. Nadir observations yield higher spatial resolution over land, while glint measurements have greater sensitivity over ocean. Each month, OCO-2 returns ~5.5 million soundings that are sufficiently cloud-free to yield full-column estimates of XCO₂ with single-sounding precisions near 0.5 ppm and accuracies <1 ppm at solar zenith angles as large as 70° (Eldering et al. 2017).

The OCO-3 spectrometer was installed on the International Space Station (ISS) in May 2019 for a planned 3-year mission lifetime. While OCO-2 collects global XCO₂ and SIF measurements from a 1:30 p.m. sun-synchronous, near-polar orbit, OCO-3 collects these

observations from dawn to dusk from the low-inclination (51°) ISS orbit. OCO-3 also includes a fast, two-axis pointing mirror assembly (PMA) enabling nadir, glint and target observations like those from OCO-2. Using the PMA, OCO-3 can change observing mode frequently within each orbit, as opposed to OCO-2 which flies in either nadir or glint mode for the entire orbit. In addition, the OCO-3 PMA enables a new observing mode called Snapshot Area Map (SAM) that returns XCO₂ and SIF measurements over targets as large as 80 km × 80 km. By combining OCO-2 and OCO-3 observations, we expect to gain a better understanding of both diurnal variations in CO₂ fluxes and urban emissions.

2.1 Instrument Characteristics

The OCO-2 and OCO-3 missions carry nearly identical spectrometers, designed to measure reflected sunlight in the O₂ A-band (ABO2) at 0.765 μm and in the weak and strong CO₂ bands at 1.61 μm (WCO2) and 2.06 μm (SCO2) with an unprecedented combination of spatial resolution, spectral resolving power, sensitivity, and dynamic range (Table 2-1 and Table 2-2).

Table 2-1. Summary of OCO-2 instrument performance.

Channel	ABO2	WCO2	SCO2
Spectral Range (μm)	0.757564 to 0.772625	1.590429 to 1.621852	2.042851 to 2.083442
Resolving Power	17961–18830	18379–20934	19150– 0420
Spectral Sampling (px)	2.46–3.37	2.34–3.03	2.25–3.03
Dynamic Range (Minimum to maximum measurable signal) (photons m ⁻² μm ⁻¹ sr ⁻¹ s ⁻¹)	7.5 × 10 ¹⁶ to 7.0 × 10 ²⁰	2.15 × 10 ¹⁶ to 2.45 × 10 ²⁰	2.15 × 10 ¹⁶ to 1.25 × 10 ²⁰

Table 2-2. Summary of OCO-3 instrument performance.

Channel	ABO2	WCO2	SCO2
Spectral Range (μm)	0.757647 to 0.772663	1.591232 to 1.622654	2.042009 to 2.082781
Resolving Power	18662–19374	19887–21237	19174–20945
Spectral Sampling (px)	2.36–3.30	2.23–3.01	2.26–3.01
Dynamic Range (Minimum to maximum measurable signal) (photons m ⁻² μm ⁻¹ sr ⁻¹ s ⁻¹)	7.5 × 10 ¹⁶ to 7.0 × 10 ²⁰	2.15 × 10 ¹⁶ to 2.45 × 10 ²⁰	2.15 × 10 ¹⁶ to 1.25 × 10 ²⁰

2.1.1 Spectral and Geometric Characteristics

The three spectral channels incorporated into each spectrometer use similar optical designs and are integrated into a common structure to improve system rigidity and thermal stability. They share a common housing and a common f/1.8 Cassegrain telescope. For OCO-2, the telescope views the Earth directly and its FOV is pointed by reorienting the spacecraft using reaction wheels. The OCO-3 telescope was redesigned with a reduced focal length and aperture to compensate for the lower altitude of the ISS and maintain footprint sizes similar to those for OCO-2. The OCO-3 telescope views the Earth through a wire-grid linear polarizer and a 2-axis (altitude, azimuth), 4-mirror (polarization compensating), pointing mechanism assembly (PMA). Because the OCO-3 field of view is not fixed relative to the platform's attitude control system, it carries a pair of context cameras to facilitate on-orbit pointing and geolocation validation.

The light path for a single spectrometer channel is illustrated in Figure 2-1. Light entering the telescope is focused on a field stop and then re-collimated before entering a relay optics assembly. From there it is directed to one of the three spectrometers by a dichroic beam splitter and transmitted through a narrow band pre-disperser filter. The pre-disperser filter for each spectral range transmits light with wavelengths within $\pm 1\%$ of the central wavelength for the band of interest and rejects the rest. The light is then refocused on the spectrometer slits by a reverse Newtonian telescope.




Figure 2-1. The OCO-2/OCO-3 instrument showing the major optical components and optical path. The figure does not include the OCO-3 PMA.

Each OCO-2 spectrometer slit is ~ 3 mm long and ~ 28 μm wide. These long, narrow slits are aligned to produce co-boresighted FOV. For OCO-2, the instantaneous FOV imaged by each slit is ~ 0.00014 radians (0.008°) wide by ~ 0.0146 radians (0.83°) long. For OCO-3, the instantaneous angular FOV is 2.2 times larger in both dimensions.

Once the light enters a spectrometer slit, it is collimated by a two-element refractive collimator, dispersed by a reflective planar holographic diffraction grating, and then focused by a two-element camera lens onto a two-dimensional focal plane array after traversing a second narrowband filter that is cooled to ~ 180 K to reject thermal emission from the spectrometer optical bench.

The spectral range of each channel is wide enough to include the complete molecular absorption band as well as some nearby continuum to provide constraints on the optical properties of the surface and aerosols, as well as absorbing gases. To meet these requirements, the O_2 A band channel covers 0.758 to 0.772 μm with a resolving power $> 17,000$, while the 1.61 and 2.06 μm CO_2 channels cover 1.591 to 1.618 μm and 2.043 to 2.083 μm , respectively, with a resolving power of $\sim 19,000$.

The spectra produced by the spectrometers are recorded by Teledyne HAWAII 1RGTM FPAs. These are 1024×1024 element FPAs with $18 \mu\text{m} \times 18 \mu\text{m}$ pixels that have a 100% fill factor

(i.e., there are no spatial or spectral gaps between pixels). The quantum efficiency of the FPAs is between 75 and 90%, and the read noise is < 30 electrons/pixel/exposure. The columns of the detector record different wavelengths, while rows record different field angles.

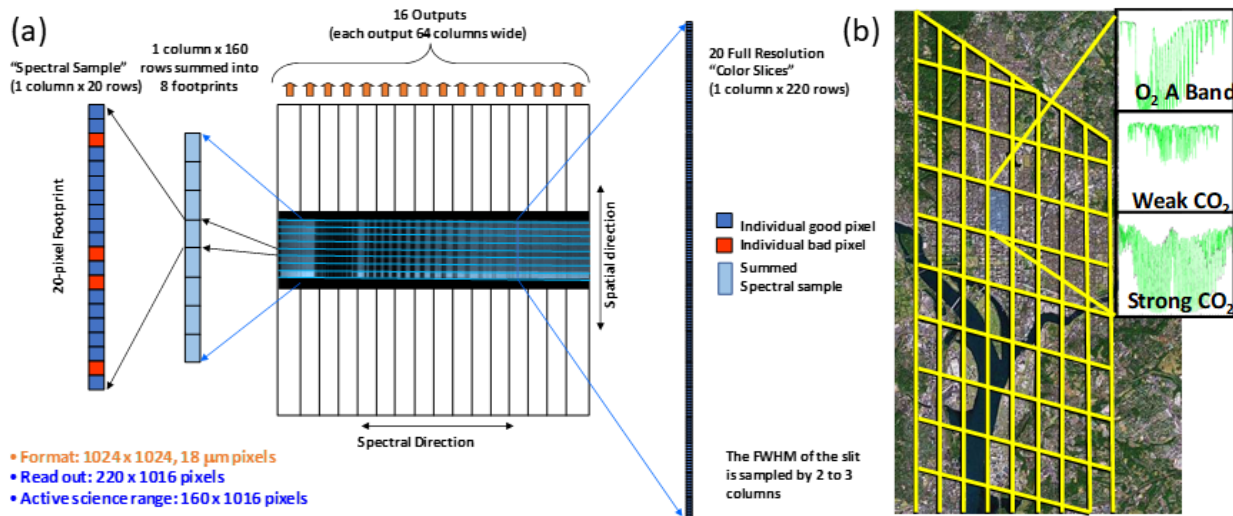


Figure 2-2. (a) The illumination and readout approach used for the OCO-2 and OCO-3 FPAs. The slits illuminate a ~ 220 (spatial) by 1016 (spectral) area of each FPA. The middle 160 pixels within the spatial field of view are summed into eight 20-pixel spatial footprints to yield eight "summed spectra" that are returned 3 times each second. (b) Spatial layout of eight cross track footprints for nadir observations over Washington, D.C. The rolling readout scheme for the FPAs, which scans over the spatial field of view, combined with the spacecraft motion, which scans along the orbit track, yields spatial footprints with a parallelogram shape. Simultaneous observations in the ABO₂, WCO₂, and SCO₂ channels are collected in each parallelogram-shaped spatial footprint and combined to form a sounding.

For OCO-2, the ABO₂, WCO₂ and SCO₂ FPA temperatures are maintained at -157 , -153 , and $-152 \pm 0.4^\circ\text{C}$, by a pulse-tube cryocooler that is thermally coupled to an external radiator through variable conductance heat pipes. The OCO-3 FPAs are maintained at temperatures about 3°C warmer by a pulse-tube cryocooler coupled to the Japanese Equipment Module Exposed Facility (JEM-EF) Active Thermal Control System (ATCS) on the ISS. At these temperatures, thermal noise from the FPAs is negligible during the short (0.333 seconds) exposure time.

For OCO-2, the optical bench Assembly (OBA) is maintained at near $-6.4 \pm 0.2^\circ\text{C}$ by a thermal radiative shroud that is coupled to an external passive radiator by variable conductance heat pipes. For OCO-3, the OBA is maintained at $-6.9 \pm 0.5^\circ\text{C}$ by a thermal radiative shroud that is coupled to a thermoelectric cooler that transfers its heat to the JEM-EF ATCS. For both instruments, the long time constant of the OBA thermal system limits temperature changes over orbital time scales, which simplifies the instrument calibration.

The spectra produced by each channel are dispersed to illuminate all 1016 pixels in the spectral dimension (columns) on each FPA (Crisp et al. 2017). The full width at half maximum of the slit image on the FPA is recorded by two to three pixels. The length of the slit limits the field of view to only ~ 190 pixels in the spatial dimension (Figure 2-2a), but 220 pixels along each FPA column are read out to record some unilluminated pixels at each end of the slit. Science observations use an along-slit field of view as defined by ~ 160 of the illuminated pixels.

The FPAs are continuously read out at 3 Hz using a “rolling readout” scheme that scans over the spatial field of view, precluding the need for a mechanical shutter and/or gaps between adjacent frames. To reduce downlink data rate and increase signal to noise ratio (SNR), ~20 adjacent pixels along an FPA column (i.e., parallel to the long axis of the slit, in the spatial direction in Figure 2-2a) are summed by the on-board processor to produce eight spatially-averaged, 1016-element spectral “footprints” across the FOV (Figure 2-2b). For OCO-2, the along-slit angular field of view of each of these spatially-averaged spectral samples (Figure 2-2a), is ~1.8 mrad (0.1° or ~1.3 km at nadir from a 705-km orbit). The angular width of the narrow dimension of the slit is only 0.14 mrad, but the telescope focus is purposely blurred to increase the effective full width at half-maximum of each slit to ~0.6 mrad. This blurring largely reduces illumination variations across the slit, which could affect the instrument line shape. It also reduces the impact of small errors in the boresight co-alignment among the three spectrometer slits. As noted above, the angular width of the OCO-3 FOV is 2.2 times larger than that of OCO-2 to maintain a similar ground footprint size. However, the OCO-3 telescope does not include as much blurring and produces substantially sharper images than those from OCO-2.

In addition to the eight spatially-binned, 1016-element spectra, each spectrometer also returns 20 columns without on-board spatial binning to provide the full along-slit spatial resolution. Each of these full-resolution “color slices” covers a 220-pixel-wide region of the FPA. This includes all of the pixels illuminated by the slit (190 pixels) as well as a few pixels beyond the ends of the slit (Figure 2-2). These full-spatial-resolution color slices are used to characterize spatial variability within each of the spatially-summed super-pixels and to monitor the thermal emission and scattered light within the instrument.

2.1.2 Polarization Characteristics

Because the diffraction gratings used in OCO-2 and OCO-3 efficiently disperse light that is polarized in the direction perpendicular to the slit, a linear polarizer is included in front of the slit. This polarizer rejects the unwanted polarization before it enters the spectrometer, where it could contribute to the scattered light background, or introduce wavelength-dependent polarization response.

If the reflected sunlight is strongly polarized and the angle of polarization of the incoming light is orthogonal to the angle of polarization diffracted most efficiently by the gratings, the signal can be substantially reduced. This issue was recognized early in the design of the original OCO instrument (i.e., Crisp et al. 2008). There was also a concern that spatially dependent changes in the angle and degree of polarization of the incoming light could introduce spectrally dependent variations in recorded spectra that could introduce regional-scale biases in the XCO₂ products. The observing strategy used by OCO-2 has been adjusted during the mission to ensure the signal levels are sufficient for the L2 retrievals, and to minimize these potential biases. OCO-3’s PMA position controls the orientation of the slit and therefore the polarization angle. Further detail on the polarization characterization is given in Section 6 of this document, and the observation strategies are described in detail in the L2 Full Physics Retrieval Algorithm ATBD.

2.2 The OCO-2 and OCO-3 Onboard Calibration Systems

Both OCO-2 and OCO-3 include on-board calibration systems. For OCO-2, the onboard calibrator (OBC) is integrated into the telescope baffle assembly (Figure 2-3). The OBC consists of a calibration paddle that carries an aperture cover (lens cap) on one end and a transmission

diffuser for solar calibration observations on the other. The backside of the aperture cover, which faces into the instrument aperture, carries a reflecting gold diffuser that can be illuminated by one of three tungsten halogen lamps.

The OCO-3 on-board calibration system uses the PMA to direct the FOV into a dedicated calibration assembly that includes three tungsten halogen lamps and a reflective gold diffuser. The PMA can also be used to observe the Moon to acquire lunar radiometric and geometric calibration measurements, but unlike OCO-2, OCO-3 cannot observe the Sun due to viewing constraints imposed by the ISS.

The OCO-2 calibration actuator rotates the paddle 90° from the “open” position to cover the telescope aperture to protect the instrument from external contamination or to monitor the dark response of the FPAs and thermal emission from the OBA. The lamps illuminate the diffuser on the back of the cover to produce spatially and spectrally smooth “flat fields” for monitoring pixel-to-pixel gain variations on the FPAs. Lamp 1 is used for routine calibration (every orbit not used for downlink), while lamps 2 and 3 are used less frequently (monthly, twice yearly) to track degradation in lamp 1, and to serve as backups if needed.




Figure 2-3. (Left) OCO-2’s Onboard Calibrator (OBC) is integrated into the telescope baffle and can be rotated to place diffusers in front of the aperture. (Right) A drawing of the OBC showing the aperture cover, solar diffuser, and the calibration paddle in the open, closed, and solar diffuser positions.

The OCO-2 calibration actuator rotates the paddle 180° from the closed position to place the transmission diffuser in front of the telescope aperture to acquire observations of Sun for radiometric and spectroscopic calibration. The solar diffuser is an all-reflective design that consists of a pair of plates separated by a cavity that is a few millimeters wide. The surfaces facing the target and the telescope aperture include arrays of pinholes that are offset to prevent direct sunlight from entering the instrument. The patterned inner surfaces of the two plates are coated with textured gold. Sunlight that enters the pinholes on the target (Sun)-facing side of the diffuser is reflected multiple times by the rough gold surfaces between the plates before reaching a pinhole on the telescope-facing side of the diffuser, where it enters the optical path to the spectrometers. The number and size of the pinholes were designed to yield diffuse intensities

similar to those from a moderately dark (< 5%) reflecting Lambertian surface when the telescope is pointed at the Sun.

The OCO-2 solar diffuser is used to acquire routine observations of the Sun just after the spacecraft crosses the northern terminator on all orbits except those that include downlinks. These measurements are used to monitor the relative radiometric response of the instrument, provided that degradation of the solar diffuser is otherwise properly accounted for. Solar Fraunhofer lines in these spectra are also used to monitor changes in instrument dispersion and changes in the instrument line shape. Early in the OCO-2 mission, these solar calibration observations were obtained while pointing the instrument field of view within $\sim 1^\circ$ of the center of the solar disk. After November 2017, the pointing was changed to offset the center of the instrument field of view 2.5° from the center of the solar disk, so that no direct sunlight would enter the instrument even if the solar diffuser was not completely covering the aperture. This change introduced a $\sim 1.5\%$ change in the throughput of the diffuser, determined by comparing measurements with and without off-pointing.

Prior to July 2019, the solar diffuser was also used periodically to acquire spectra of the Sun over the entire illuminated part of an orbit. These observations sample the full range of solar relative velocities (± 7 km/s) and associated Doppler shifts observed over the illuminated hemisphere. The Doppler-shifted spectra are combined to produce a massively over-sampled solar spectrum, which is used to monitor subtle variations in the ILS. These were also used to refine understanding of the top-of-atmosphere solar spectrum for the wavelength ranges sampled by OCO-2. Solar Doppler observations were suspended after 14 June 2019 due to the failure of the OCO-2 inertial measurement unit, because the Earth's disk blocks the star tracker's field of view during these observations.

The OCO-2 calibration paddle is rotated 90° from either the closed or diffuser positions to expose the telescope aperture for routine science and lunar calibration observations. OCO-2 acquires lunar observations to monitor the instrument's geometric calibration and to monitor the degradation of the instrument without a solar diffuser. Prior to July 2019, lunar calibration observations of the 75% illuminated waxing gibbous and near-full phases of the Moon were acquired each month. Full moon observations were suspended after 14 June 2019 due to the failure of the OCO-2 inertial measurement unit, because the Earth's disk blocks the star tracker's field of view during these observations. The Mission Operations Team is exploring which other Lunar phases can be safely observed without blocking the star tracker.

OCO-3 can use its PMA to collect lunar calibration observations as the lunar disk moves through a small field of view between the Earth and the ISS. Unlike OCO-2, which scans across the lunar disk only once during each lunar calibration, the OCO-3 PMA can scan across the lunar disk 2–6 times in a given orbit. However, from its position on the ISS, it cannot execute these observations every lunar month due to obstructions in its field of view. Because OCO-3 cannot observe the Sun and observes the Moon less frequently than OCO-2, observations of well calibrated surface vicarious calibration sites and cross calibrations of coincident observations over pseudo invariant calibration sites (PICS) will play a more significant role in future versions of OCO-3 on-orbit radiometric calibration.

2.3 Types of On-Orbit Calibration Measurements

The radiometric, spectroscopic, and geometric calibration requirements for the OCO-2 and OCO-3 instruments are particularly demanding because their Level 2 products must resolve changes smaller than 0.25% in the column-averaged CO₂ dry air mole fraction XCO₂. For example, absolute radiometric errors greater than 5% or pixel-to-pixel relative radiometric errors larger than a few tenths of one percent can introduce biases in the Level 2 XCO₂ products larger than 0.25%.

To meet these demanding calibration requirements, the OCO-2 and OCO-3 instruments were extensively characterized and calibrated prior to launch (Frankenberg et al. 2015; Rosenberg et al. 2017; Lee et al. 2017). Once in orbit, both routine and special calibration observations are collected to detect and correct changes in instrument response over time. For example, the instrument radiometric throughput can change in response to the accumulation of contaminants on the spectrometer optics or focal plane arrays (FPAs), changes in the optical bench or FPA temperatures, degradation of optical coatings, and changes in the performance of individual pixels on the FPAs (Crisp et al. 2017; Marchetti et al. 2019). To monitor changes in the dark response or gain of the FPAs or the throughput of the optical system, OCO-2 collects dark, lamp, and solar observations on every orbit that is not used for downlink. Lunar observations are collected once or twice each month to provide independent estimates of the throughput of the science optical path and to monitor changes in the throughput of the solar diffuser. On longer time scales, observing campaigns over well-characterized surface targets, such as Railroad Valley, Nevada (Bruegge et al. 2019) and pseudo invariant calibration sites (Yu et al. 2020) help to detect and correct long-term drifts in radiometric calibration. The types of data collected to monitor the stability of the instrument calibration are summarized in Table 2-3.

Table 2-3. Summary of inflight calibration measurements.

Calibration Type	Approach, Purpose, and Results
Solar Calibration (only for OCO-2)	<p><u>Approach:</u> View the Sun through the transmission diffuser when in the vicinity of the northern terminator (shortly after final science measurements for a given orbit)</p> <p><u>Frequency:</u> ~6x/day each in single pixel and summed modes</p> <p><u>Purpose:</u> Relative radiometric and spectroscopic trending</p> <p><u>Results:</u> Critical for tracking radiometric degradation, ages more slowly than lamps and has stable spectral slope.</p>
Lunar Calibration	<p><u>Approach:</u> Point the instrument boresight at a right ascension/declination location and allow the spacecraft orbital motion to scan the lunar disk</p> <p><u>Frequency:</u> for OCO-2, ~1x per lunar cycle at ~75% waxing phase and 1x per lunar cycle near full phase prior to 14 June 2019.</p> <p><u>Frequency:</u> For OCO-3, as often as possible given ISS viewing constraints, roughly 5 opportunities per year to observe several times over three days</p> <p><u>Purpose:</u></p> <ul style="list-style-type: none"> • Provide pointing validation • Provide relative radiometric calibration (to trend other calibration data, e.g., solar calibrator degradation) <p><u>Results:</u> OCO-2 has acquired over 100 observations which are vital to understanding radiometric degradation without an aging lamp or aging solar diffuser. OCO-3 lunar results from 2020–2022 have helped to understand pointing, but not radiometric trends.</p>

Calibration Type	Approach, Purpose, and Results
Lamp Calibration	<p><u>Approach:</u> View one of three onboard lamps using the internal diffuser</p> <p><u>Frequency:</u> for OCO-2</p> <ul style="list-style-type: none"> • Lamp #1: ~6x/day each in single pixel and summed modes • Lamp #2: ~12x/year in both single pixel and summed mode • Lamp #3: ~2x/year in both single pixel and summed mode <p><u>Frequency:</u> for OCO-3</p> <ul style="list-style-type: none"> • Lamp #3: 80/100 orbits (~12x/day) each in single pixel and summed modes • Lamp #2: 19/100 orbits (~3x/day) in both single pixel and summed mode • Lamp #1: 1/100 orbits (~ 1x/week) in both single pixel and summed mode <p><u>Purpose:</u> Provide relative radiance standard with no spectral lines</p> <p><u>Results:</u> Slow reduction in intensity (lamp output, color temperature, and diffuser reflectance) and color temperature thought to be due to lamp aging, but provides means to track pixel-to-pixel radiometric changes.</p>
Dark Calibration	<p><u>Approach:</u> Collect data with lamps off and internal diffuser blocking aperture</p> <p><u>Frequency:</u> ~20x/day for OCO-2 each in single pixel and summed modes, ~40x for OCO-3</p> <p><u>Purpose:</u> Monitor fixed pattern changes and fit out optics and focal plane temperature dependencies</p> <p><u>Results:</u> Dark subtraction residuals no more than 20% background noise of instrument</p>
Full Orbit Darks (only for OCO-2)	<p><u>Approach:</u> Collect data with lamps off and internal diffuser blocking aperture during the day side of the orbit</p> <p><u>Frequency:</u> During in orbit checkout, after that only as needed such as after a decontamination</p> <p><u>Purpose:</u> Validate that dark model developed with data taken in eclipse is useful for correcting science data offsets</p> <p><u>Results:</u> Data on day side is statistically similar to dark data taken in eclipse</p>
Solar Doppler Diffuse (OCO-2 prior to 14 June 2019 only)	<p><u>Approach:</u> Maintain solar calibration during entire dayside of an orbit</p> <p><u>Frequency:</u> ~1x/month or after each decontamination cycle (note that when this occurs, an orbit of science data is lost)</p> <p><u>Purpose:</u> Take advantage of $\sim\pm 5$ km/s Doppler shift of the solar lines to validate the instrument line shape stability</p> <p><u>Results:</u> Full width at half-maximum of ILS stability to <1% (repeatability of measurement)</p>
Vicarious Calibration	<p><u>Approach:</u> Track Railroad Valley, NV as spacecraft flies over collecting 1000s of spectra within a few km of the center of the playa, then compare to field measurements of surface reflectance</p> <p><u>Frequency:</u> ~10x/year but frequency is driven by weather and field campaigns to collect simultaneous data at the site (note that when this occurs, ~1/4 orbit of science data is lost)</p> <p><u>Purpose:</u> Validate absolute calibration of radiometric output</p> <p><u>Results:</u> Differences between OCO-2 and ground measurements <4% for observations acquired on orbit path 138 (Bruegge et al. 2019).</p>
Cross Calibration	<p><u>Approach:</u> Utilize nadir data collected over invariant sites (primarily Sahara Desert) to create a timeseries intercomparison to other sensors (MODIS, VIIRS etc.) for trend analysis.</p> <p><u>Frequency:</u> Up to once every 16 days per site. OCO-2 and Aqua MODIS had 3–8 usable overpasses per year at each of eight different PICS sites. OCO-2 and OCO-3 also have dozens of simultaneous nadir overpasses per month, though many are over ocean.</p> <p><u>Purpose:</u> Assess long-term radiometric calibration stability in the L1b data product</p> <p><u>Results:</u> For OCO-2 Version 8 (v8r) product, quantified long-term stabilities are consistent with the ones derived using the lunar observations (Yu et al. 2020)</p>

Calibration Type	Approach, Purpose, and Results
Streak Flats (OCO-2 Only)	<p><u>Approach</u>: The OCO-2 spacecraft rotates to maintain polarization, leading to a few seconds each orbit where the slit orientation is parallel to the ground track and footprints from consecutive frames overlap.</p> <p><u>Frequency</u>: Once per orbit</p> <p><u>Purpose</u>: Assess radiometric differences between footprints</p> <p><u>Results</u>: Spectrally flat adjustments to preflight gains of up to 2%</p>

Changes in the instrument thermal environment or other factors that change the focus or alignment of the spectrometers can introduce changes in the spectral dispersion or instrument line shape, compromising the spectral calibration (Crisp et al. 2017). Wavelength shifts and linear stretches in the dispersion in each channel are monitored by measuring the positions of O₂ and CO₂ lines in the science spectra. Changes in the instrument line shape are tracked using observations of the positions and widths of solar Fraunhofer lines in the solar spectra collected during routine solar calibration and solar Doppler experiments. Zero-level offsets (ZLO) due to changes in the internal scattering are characterized by monitoring regions of the FPAs that are read out in a diagnostic “single pixel” calibration mode but not directly illuminated by the spectrometer slits.

Small geolocation errors can also introduce XCO₂ uncertainties much larger than 0.25% by introducing uncertainties in the optical path length, particularly over land areas with significant topographic variability. For OCO-2, the preflight requirement for pointing accuracy was one milliradian (1 mrad = 0.057°). This requirement was validated during the in-orbit check out period using observations of the lunar disk and coastline crossings. However, analysis of V8 data products revealed significant (± 3 ppm) variations in XCO₂ in regions with topographic changes (Wunch et al. 2017; Kiel et al. 2019). An investigation traced these biases to a small (0.03°) pitch error in the pointing of the instrument boresight that was corrected in the OCO-2 B10 product. This more stringent requirement has also been adopted as a goal for OCO-3. The pointing and geolocation calibration approach is described in Section 5.1.

In addition to this spectrally independent geolocation calibration requirement, all spectral samples within the full spectral ranges of the three spectrometer channels must record light from the same spatial field of view, or optical pathlength differences will introduce spectrally-dependent calibration errors. For imaging grating spectrometers, this spatial-spectral requirement can be met by aligning the spectrometer slits, the grooves on the diffraction gratings, and columns of the FPAs. Each spatial location along the slit is then sampled by a constant set of rows on the FPAs. For the OCO-2 and OCO-3 instruments, this alignment is very good, but not perfect, and the FPAs are slightly rotated (or “clocked”) relative to the slit and grating. Consequently, a given spatial footprint does not map onto the same set of rows on the FPA across the entire spectral range, but instead varies (roughly linearly) with spectral position (i.e., column). To compensate for this, and record the same spatial information across the entire spectrum, the starting row index for each spectral sample can be adjusted in increments of one pixel ($\sim 1/20$ th of a summed footprint). This approach introduces little error in spatially homogeneous scenes, but can produce radiance discontinuities in spectra of scenes with strong intensity variations near the edge of a footprint. These discontinuities are corrected as part of the calibration process (see Section 5).

2.4 Data Product Delivery

Science and housekeeping data for OCO-2 are transmitted to a NASA Near Earth Network station once each day. The data are then transferred to the Earth Science Mission Operations center at the NASA Goddard Space Flight Center (GSFC), where the raw telemetry is converted to time-ordered raw radiance spectra (Level 0 Products). For OCO-3, the raw telemetry data is brought through White Sands to the Marshall Space Flight Center (MSFC). For both missions, Level 0 products are then delivered to the OCO-2 Science Data Operations System at the NASA Jet Propulsion Laboratory (JPL), they are processed to yield Level 1A and then Level 1B Products. The OCO-2 and OCO-3 Level 1B (L1B) product files usually cover the full day side of an orbit with calibrated and geographically located spectral radiances from the ABO2, WCO2, and SCO2 channels and associated metadata. Each file or “granule” included in this product contains a record for each spectral sample and footprint in each channel for each of the ~74,000 soundings the instrument collects while viewing Earth during a single spacecraft orbit. OCO-2 delivers granules containing only one mode (nadir, glint, target, transition) of data, so an orbit will have either 1 or 5 granules. OCO-3 switches modes frequently, and delivers one “science” mode granule containing many different sub-modes.

The OCO-2/OCO-3 calibration teams deliver calibration updates to the Science Data Operations System for use in generating the L1B products in the form of Ancillary Radiometric Product (ARP) and Ancillary Geometric Product (AGP) files. The ARP includes updates to the dark correction, gain degradation, ZLO, SNR, bad sample lists, declocking, ILS, and spectral dispersion offset and stretch terms. The AGP includes the footprint center locations and FOVs of all three spectral channels, relative to the instrument reference system frame which is based on the telescope chief ray. It also includes the polarization angles for each footprint in each spectral channel.

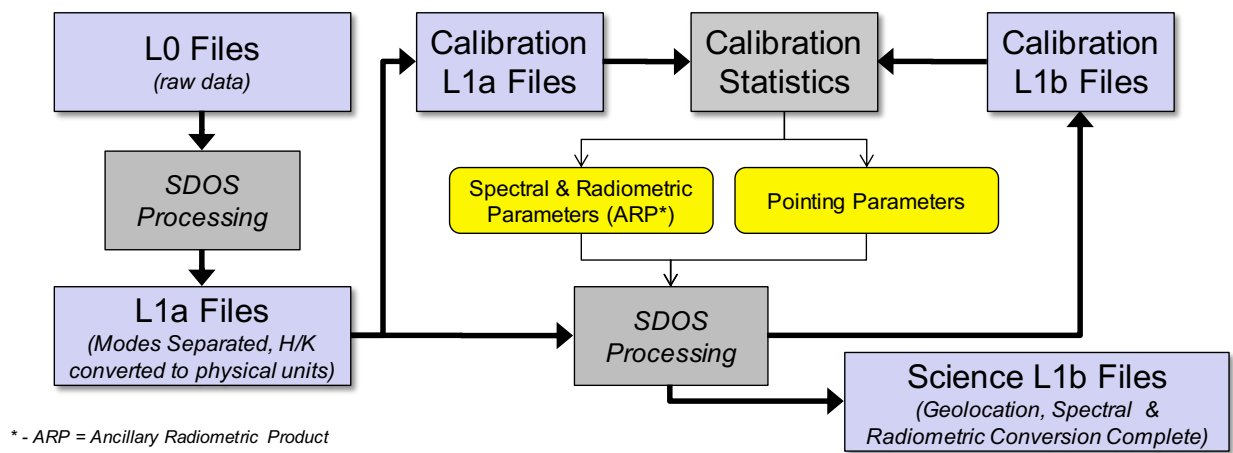


Figure 2-4. A schematic diagram illustrating the process for generating Level 1B files from Level 0 files.

Two versions of the OCO-2 and OCO-3 L1B products are routinely delivered to the Goddard Earth Science Data and Information Services Center (GES DISC). The “forward” product, designated V11, uses calibration information based on extrapolations of calibration data collected in the recent past. These products are usually delivered to the GES-DISC within a week of acquisition. The V11 products are generally reliable except when the instrument experiences a

significant change in its thermal environment that is outside the range included in the extrapolated calibration data. The “retrospective” L1B and L2 products, designated V11r, use interpolated calibration data and are therefore expected to be more reliable than the forward products. The V11r products are typically delivered to the GES-DISC in one-month blocks, starting about 2–6 weeks after the date of acquisition.

3 Radiometric Calibration

The radiometric performance and calibration of the OCO-2 and OCO-3 instruments were established prior to launch through an extensive series of measurements acquired while the instruments were in flight-like conditions in a thermo-vacuum chamber, viewing an integrating sphere with NIST-traceable calibration standards (Rosenberg et al. 2017). The onboard calibration systems were also calibrated against these references in these tests and are now being used to trend the in-flight instrument radiometric calibration.

The OCO-2 and OCO-3 radiometric calibration process is summarized in Figure 3-1. The focal planes in each spectral channel return unitless data numbers (DN). Time ordered geolocated spectra, expressed in DN, are referred to as the L1A spectra. The raw DN values include contributions from the bias and dark current and any photoelectrons collected by each pixel or spectral sample in the active area of the FPA. The first data processing step is to apply dark correction, which removes the bias and dark current from each spectral sample. As discussed in Rosenberg et al. (2017), this is important because a DN offset can introduce a linearity error. During thermal-vacuum testing, a set of data was collected with the light sources shuttered such that no light entered the instrument. These dark data were combined to produce high signal-to-noise single pixel and summed “synthetic darks.” These data are used to create a baseline for trending the instrument performance, and were used for dark corrections early in the mission.

While on orbit, dark data are routinely collected during the eclipse phase of the vast majority of orbits. The dark data collected over several orbits are combined to improve their signal to noise ratio and are used to correct the zero-illumination response of the data included in the L1B

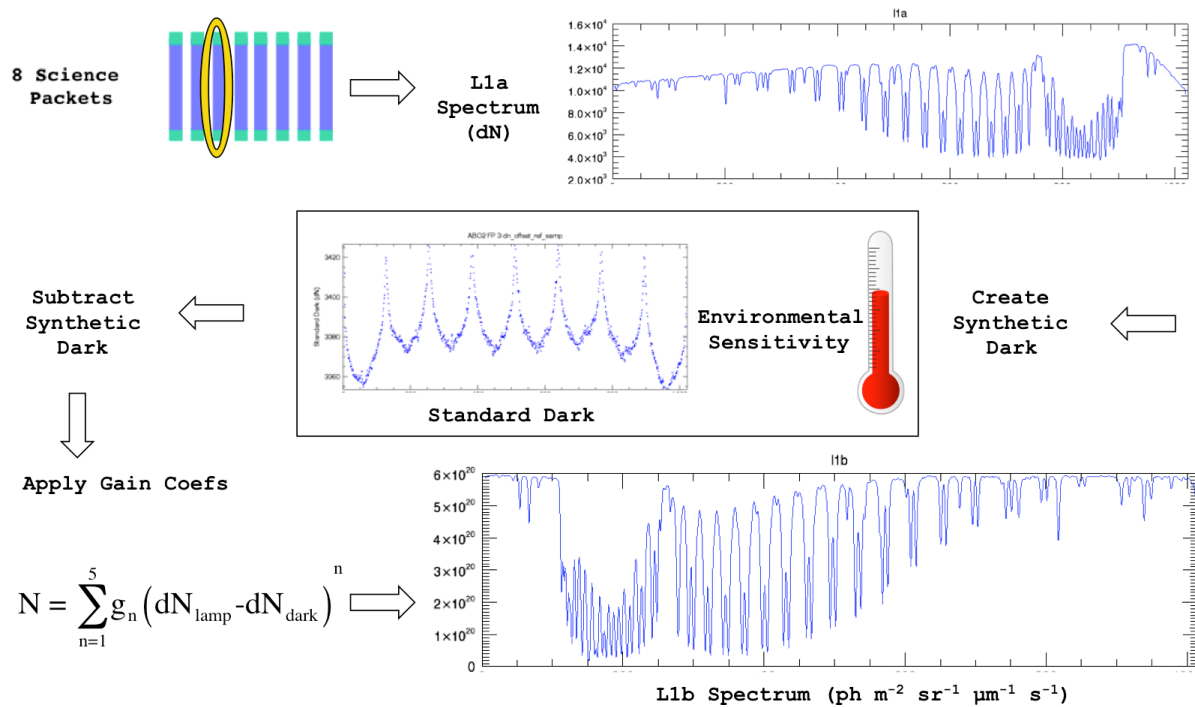


Figure 3-1. Overview of the ground processing data flow for the radiometric calibration process used to generate L1B data. Stray light corrections and declocking are omitted.

product as a function of two temperatures. Forward calibration always uses the focal plane and optical bench temperatures, while retrospective calibration chooses the pair of temperatures that explains the greatest variance in the DN data. The methods used to create and apply the in-flight synthetic darks are described in greater detail below. After dark correction, a stray light (“ZLO”) correction proportional to the science-area-average signal is applied. The next step is to apply gain curves, which describe the relationship between DN and calibrated radiance. The preflight gain curves are described in more detail in Rosenberg et al. (2017).

A new set of calibration coefficients are derived at roughly weekly intervals, integrated into an Ancillary Radiometric Parameter (ARP) file and applied to approximately one week of data. This weekly cadence has proven to be a good balance between the number of calibration observations needed to reduce the noise in the various curve fits, and capturing transitions in instrument behavior. The precise length of each period varies, with boundaries determined by operational events such as thermal setpoint adjustments. The sections below describe the approach for inflight updates.

3.1 Bad Pixels and Bad Samples

OCO-2 is flying OCO flight-spare Teledyne Hawaii 1RG HgCdTe FPAs, which were delivered in 2006. A few percent of the 225,280 pixels in the active area of these arrays are either “dead”, or respond to light or thermal changes in a way that is not consistent with the majority of the other pixels. These pixels must be identified and excluded from the 20-pixel sums that are performed onboard, or they will contaminate the resulting 20-pixel “spectral samples” that constitute the spectra returned for the eight footprints by each spectrometer.

Since bad pixels are removed before downlink, no additional bad pixel corrections can be made to data that has already been acquired. When new outliers appear, they are marked for a future update. The only way to mitigate their effect on the science data is to remove the ~20-pixel spectral sample. Spectral samples with too many bad pixels, or with other known issues (e.g., optical artifacts on the short-wavelength ends of all three bands) are marked as bad in the L1B product, and should only be used with caution (almost all users should ignore them in higher-level processing). Because eliminating bad pixels within a spectral sample can change the sample’s radiometric bias and gain, the upload of a new bad pixel map also requires changes to several calibration coefficients, and therefore forms a major boundary in a reprocessing campaign.

The bad pixel identification process started during prelaunch testing, but must be updated on orbit because additional pixels degraded during the 2.5-year instrument storage period prior to launch, and others are damaged on-orbit by cosmic radiation, thermal cycling of the FPAs, and other effects. To identify bad pixels, the calibration team routinely collects dark and lamp data using “single pixel” mode, which returns all of the pixels in the active region of each array. (Single-pixel results cannot be returned routinely for science observations because it requires ~9 seconds to read out an FPA at full resolution and therefore would lead to a ~97% reduction in obtained spectra.) These single-pixel calibration data are then analyzed to identify bad pixels and add them to a “bad pixel map” that is uploaded to the spacecraft for use in the pixel-summing process. This has been done five times since launch on OCO-2 (see Table 3-1 and once on OCO-3 (see Table 3-2). As the bad pixel maps improve, a larger fraction of the spectral samples in each spectrum can be used to generate L2 products.

Table 3-1. History of OCO-2 bad pixel map uploads. The bad pixel map version numbers show the three spectral channels (ABO2/WCO2/SCO2). The ABO2 channel has fewer bad pixels and has been updated less frequently than the WCO2 and SCO2 channels.

Bad Pixel Map Version #	Date Loaded	Notes
5 / 5 / 5	preflight	Derived during preflight tests in 2012 and used during in-orbit checkout. No science data are being delivered from this period.
5 / 6 / 6	2014-09-05T20:06:06.290Z	These updates to the CO ₂ band bad pixel maps removed pixels that had lost sensitivity to light between initial characterization and launch.
5 / 7 / 7	N/A	Tested but never uploaded.
5 / 8 / 8	2014-11-10T14:56:20.120Z	First updates to CO ₂ bad pixel maps to mask out pixels with very high noise or very unstable dark current.
5 / 9 / 9	2014-11-21T16:16:54.188Z	Improvement to 5/8/8 based on initial review of data
5 / 10 / 10	2015-02-13T15:53:08.785Z	Further refinements to remove additional noisy/unstable pixels.
11/11/11	2018-08-09T12:28:52.142Z	First bad pixel map created with the machine learning classifier (Marchetti et al. 2019)

Table 3-2. History of OCO-3 bad pixel map uploads. The bad pixel map version numbers show the three spectral channels (ABO2/WCO2/SCO2).

Bad Pixel Map Version #	Date Loaded	Notes
102/102/102	preflight	Derived during preflight tests in 2018 and used during in-orbit checkout. No science data are being delivered from this period.
103/103/103	2019-07-26T01:11:26.903Z	Combined new bad pixels from the end of preflight testing and the first month in space.

The process for identifying bad pixels was significantly improved in 2018 with the introduction of a machine learning classifier. The classifier used the same lamp and dark data as before, but defined several features such as similarity to neighbors and change over time. The features, along with a prior bad pixel map, are combined to give a likelihood that a given pixel is bad. The calibration team then chooses a threshold likelihood to determine the final map. The classifier is run each two months for each mission, and has indicated very stable performance from 2019–2022.

OCO-3 uses a flight spare Teledyne Hawaii 1RG HyViSi from the 2006 OCO lot, but it uses a new lot of HgCdTe FPAs for the WCO2 and SCO2 channels. These newer HgCdTe FPAs have many fewer bad pixels than those used by OCO-2. An initial bad pixel map was determined manually in 2013, then updated twice during preflight testing and once during in-orbit checkout using the ML classifier.

3.2 Bad Pixel Map

The application of the bad pixel map is invisible to users of the standard L1B product. The following is presented for information only.

The bad pixel map is a $1024 \times 220 \times 3$ byte array where a value of 1 indicates that a pixel's signal should not be used. Bad pixels are eliminated by the flight computer in space, while producing sample mode (“summed”) data. The full single pixel readings are not routinely returned to the ground during nominal science operations, but summed mode products do the color slices at 3 Hz.

Within each ~20-pixel spectral sample, a single or double bad pixel is replaced by the average of the closest neighboring good pixels on each side. This is a good approximation because the spatial imaging of the telescope has been blurred to yield a point spread function with a full width of 3–6 pixels. If there are three or more contiguous bad pixels within a single spectral sample, those values are set to 0 and not included in the sum. As a result, a pixel can have a weight of 0, 1, 1.5, 2, 2.5, or 3 in the nominal 20 pixels that constitute a sample. Many sample mode calibration fields (gain, noise model, dark correction coefficients) must be updated with each bad pixel map change to reflect the new weighting.

3.2.1 Bad Pixel Map Usage

The bad pixel map is used only in SUMMED mode. The bad pixel map is not applied to data collected in single pixel mode. The bad pixel map also is not applied to the color slices. These data must be corrected separately by the user.

3.2.2 Summed Mode Pixel Map Processing

All bad pixels identified in a bad pixel map are replaced, as described above, before the onboard instrument software applies any pixel summation algorithms.

3.2.3 Bad Pixel Map/Footprint Relationship

The OCO-2 instrument digitizes an array of 1024 spectral pixels (columns) by 220 spatial pixels (rows). Summed science data uses the “middle” 160 spatial pixels. This relationship was described in Section 2.1.1 and is shown in Figure 2-2. The eight footprints are each approximately twenty pixels wide, with footprint limits that vary by band as shown in Table 3-3.

The pixel map is always referenced to the 1024×220 digitized array but in order to be most useful, the pixel map limits should correspond to the limits defined in the footprint start and end limits, plus two rows on each end. Thus, using the example, the smallest ABO2 row containing bad pixels should be 31, and the largest row with bad pixels should be 192.

Table 3-3. OCO-2 footprint limits at the long wavelength end.

Footprint	ABO2	WCO2	SCO2
1	33 – 51	35 – 54	26 – 45
2	52 – 71	55 – 73	46 – 65
3	72 – 91	74 – 93	66 – 85
4	92 – 111	94 – 113	86 – 105
5	112 – 131	114 – 133	106 – 125
6	132 – 151	134 – 153	126 – 145
7	152 – 171	154 – 172	146 – 165
8	172 – 190	173 – 192	166 – 185

3.2.4 Bad Pixel Mitigation Calculations

Bad pixels are corrected as shown in Table 3-4. Bad pixel correction is only performed in the spatial direction within the same column. These corrections are performed prior to the footprint summing process, without regard to footprint limits or boundaries. No interpolation or averaging is performed in the spectral direction.

Table 3-4. Bad pixel mitigation.

Bad Pixel Configuration	Mitigation
Single bad pixel	Replaced with the average of the two spatially adjacent good pixel values
Double bad pixel	Replaced with the average of the two spatially adjacent good pixel values; the same value is used to replace both bad pixel values
Triple bad pixel or larger	All bad pixels are replaced with zero values

3.3 Dark Subtraction

Variations in the temperature of the optics and FPAs contribute to the dark background signal. The as-measured DN values can be corrected for these effects using the dark subtraction equation:

$$dn_{corrected} = (dn_{raw} - dn_{ref}) + c_{optics} (T_{optics}(t) - T_{ref_optics}) + c_{FPA} (T_{FPA}(t) - T_{ref_FPA}) \quad (3-1)$$

where

- $dn_{corrected}$ —DN after dark correction
- dn_{raw} —as-measured DN
- dn_{ref} —DN detected with no illumination
- c_{optics} —optics temperature dependence coefficients
- $T_{optics}(t)$ —optics temperature as a function of time (in L1B data product as FrameTemperatures/temp_smooth_optical_bench_grating_mz)
- T_{ref_optics} —the optics temperature that will yield dnref
- c_{FPA} —FPA temperature dependence coefficients
- $T_{FPA}(t)$ —O₂, weak CO₂, and strong CO₂ FPA temperatures as a function of time (in L1B data product as FrameTemperatures/temp_smooth_fpa_[o2, weak_co2, strong_co2])
- T_{ref_FPA} —the O₂, weak CO₂, and strong CO₂ FPA temperatures that will yield dn_{ref}

dn_{raw} , dn_{ref} , c_{optics} , and c_{FPA} have dimensions of $1024 \times 8 \times 3$ —a value for each column, footprint, and band (Figure 3-2). These quantities are included in the Ancillary Radiometric Product (ARP) but are not provided in the L1B files.

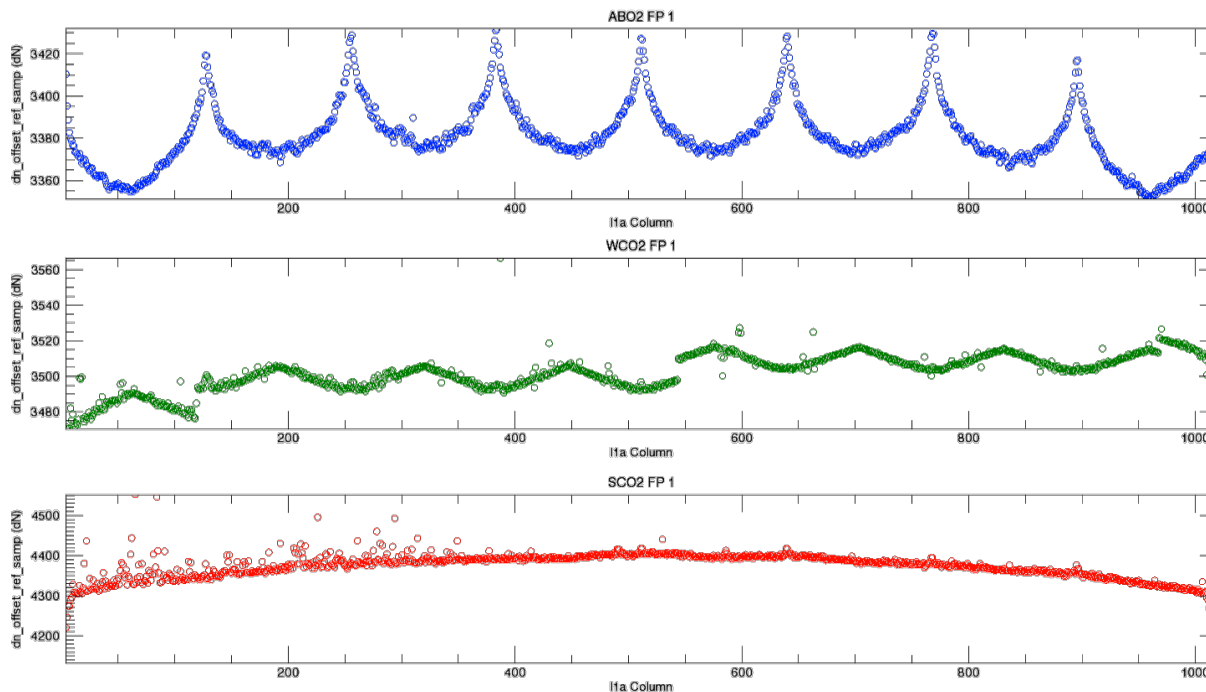


Figure 3-2. Example of a term in the OCO-2 dark subtraction (dn_{raw}). Spectral patterns in dark coefficients reflect features of the detectors (ABO2, WCO2) and optics (SCO2), and remove these artifacts from L1B spectra.

The bias and dark current for an individual pixel depends on voltage and temperature variations within the instrument. To minimize the science impact of these changes, dark corrections performed in the conversion from L1A to L1B include simultaneous adjustments for two temperatures. To avoid introducing errors due to noise in the temperature sensors, temperatures are smoothed before they are used in this correction algorithm. Note that all temperatures except for the FPA temperatures are only updated every 15 seconds for OCO-2. Temperatures are smoothed over an entire orbit using a linear fit, performed before the orbit is broken into several files at mode change boundaries. One impact of this approach is a poor temperature fit (and thus dark correction) for orbits that include adjustments to the optics or focal plane temperatures.

3.4 Radiometric Gain

Because atmospheric absorption is inferred from the ratio between continuum and line core radiances, the radiometric calibration must account for any nonlinearities in the instrument gain. Although the instrument response appears almost linear, up to a 5th order polynomial of radiance vs. signal is provided for to describe any nonlinear effects:

$$Radianc = \frac{1}{k} \sum_{i=0}^5 c_i * dn_{darkcorrected}^i \quad (3-2)$$

The scaling terms, k , are stored in the variable `gain_degrad_samp` and were set to 1.0 at launch. In-flight, the `gain_degrad_samp` terms are used to track changes in the instrument response throughout the mission through the vicarious calibration experiments described in Section 1.4. The c_i terms are stored in the variable `gain_preflight_samp`. At launch, these are set based on a combination of the best radiometric calibration tests performed during final instrument testing at JPL. While the algorithm allows an offset term ($i = 0$), c_0 is always set to zero as dark subtraction is made in the earlier computational step. No significant improvement was observed after cubic correction, so the c_4 s and c_5 s are also set to zero. For example, if $c_1 = 2.898 \times 10^{15}$, $c_2 = 1.902 \times 10^9$, $c_3 = 9.559 \times 10^3$, and $c_0 = c_4 = c_5 = 0$, then the equation above becomes:

$$\text{Radiance} = 0 + 2.898 \times 10^{15} * dn + 1.902 \times 10^9 * dn^2 + 9.559 \times 10^3 * dn^3 + 0 * dn^4 + 0 * dn^5 \quad (3-3)$$

The determination of the `gain_preflight_samp` coefficients is described in detail in Rosenberg et al. (2017). These values are being updated to minimize footprint-to-footprint variations.

3.5 Radiometric Gain Corrections

The radiometric response of the instrument has been the most variable aspect of in-flight performance and has required challenging calibration updates, which are described below.

To monitor the radiometric response of the OCO-2 spectrometers, frequent observations of the onboard lamps and the sun, along with less frequent observations of the Moon and the vicarious calibration site Railroad Valley, NV, are used. Not all of the changes in the regular calibration measurements can be attributed to instrument degradation, however, as both the output of the lamps and the reflectance of the lamp and solar diffusers change over time. Ultimately, lunar observations have proved key in determining the steady absolute degradation of the instrument, while the other measurements allow relative changes to be well characterized as a function of wavelength and time.

As described below and in more detail in Crisp et al. (2017), three phenomena affecting the radiometric response of the instrument have been identified for OCO-2. They are: a fast, reversible degradation, a slow irreversible degradation, and a zero-level offset. Lunar, lamp and solar calibration data were critical to differentiate and correct for these effects.

3.5.1 Fast Degradation

ABO2 solar observation timeseries data reveals a series of cycles of rapid and accelerating signal loss followed by nearly full signal recovery (Figure 3-4). These cycles, which begin and end on instrument decontamination events (decons), are due to accumulation of a thin layer of ice on the FPA. The ice acts to reduce the effectiveness of the AR coating, thereby reducing radiometric throughput. Because the thickness of the ice layer varies over the surface of the FPA, the resulting loss of throughput varies not only temporally but also spatially and spectrally. Ice accumulation occurs on the WCO2 and SCO2 FPAs, too, but the impact on these two channels is much smaller owing to differences in detector material and AR coating.

3.5.2 Slow Degradation

Figure 3-5 shows the solar calibration signal following each instrument decontamination cycle. To date, this gradual, irreversible loss in throughput has reached roughly 14% in the ABO2 channel versus less than 2% in the CO₂ channels. Early in the mission, the loss in solar signal

was attributed, for want of evidence to the contrary, to radiometric degradation of the instrument. Based on observations of the Moon it was subsequently concluded, however, that the majority of this signal loss was in fact due to changes in the solar diffuser. Whereas the solar cal signal does not, therefore, constitute a reliable indicator of instrument slow degradation, it is critical for characterizing fast degradation as a function of time. To this end, the irreversible loss of solar signal in between decons must be estimated from the post-decon data. Prior to B10, this was done by refitting a curve to the post-decon data following each decon. As the changing fit coefficients had an effect on data far in the past, however, since B10 linear interpolation has been used to fill in the gaps between decons.

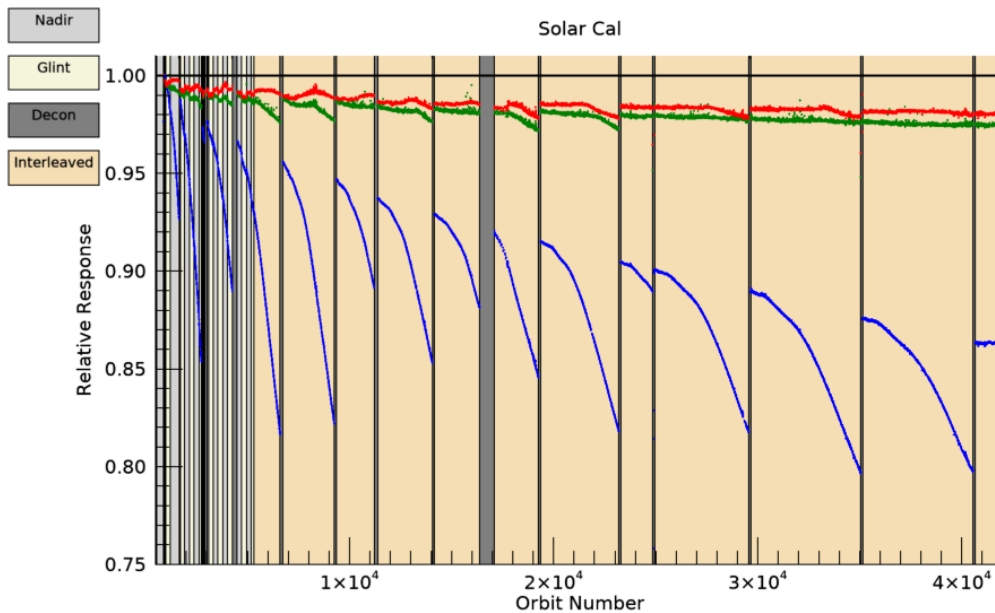


Figure 3-3. OCO-2 relative solar cal radiometric response for the ABO2 (blue), WCO2 (green), and SCO2 (red) channels as a function of time. The dark bands indicate decontamination events, during which the instrument optical bench is warmed to 12°C to remove ice from the cryolinks and FPA surfaces. The ABO2 solar signal exhibits a fast degradation process that is reversed during these decons superimposed upon a slow degradation process that is not.




Figure 3-4. Slow loss of solar calibration signal in the (a) ABO2 channel, (b) WCO2 channel, and (c) SCO2 channel. Each point is the initial band mean relative radiometric response following a decon. These time series track the combined slow degradation of the solar calibration diffuser and the instrument optics.

With both the lamps and the solar diffuser undergoing more rapid degradation than the instrument, observations of the Moon, in combination with solar cal data to correct for reversible, fast degradation, have been deemed to provide the most accurate measurements of channel-mean instrument slow degradation. In addition to correction for icing, OCO-2 lunar observations require correction for observing geometry, including as the impacts the polarization state of the light reflected from the Moon. The observing geometry-based corrections are derived from a parametric fit to the lunar cal time series. The results of applying these, plus the icing corrections to the lunar cal time series are shown in Figure 3-5. The A-Band shows an degradation of approximately 0.5%/year; in contrast, the CO₂ channels exhibit only very slight annual degradation, based on the existing set of lunar observations. The lunar-cal-derived slow degradation rates are fed into the calibration pipeline. The vicarious calibration data collected at Railroad Valley, NV, show that instrument radiometric change over time has been accurately corrected, to within the uncertainty of the measurement.

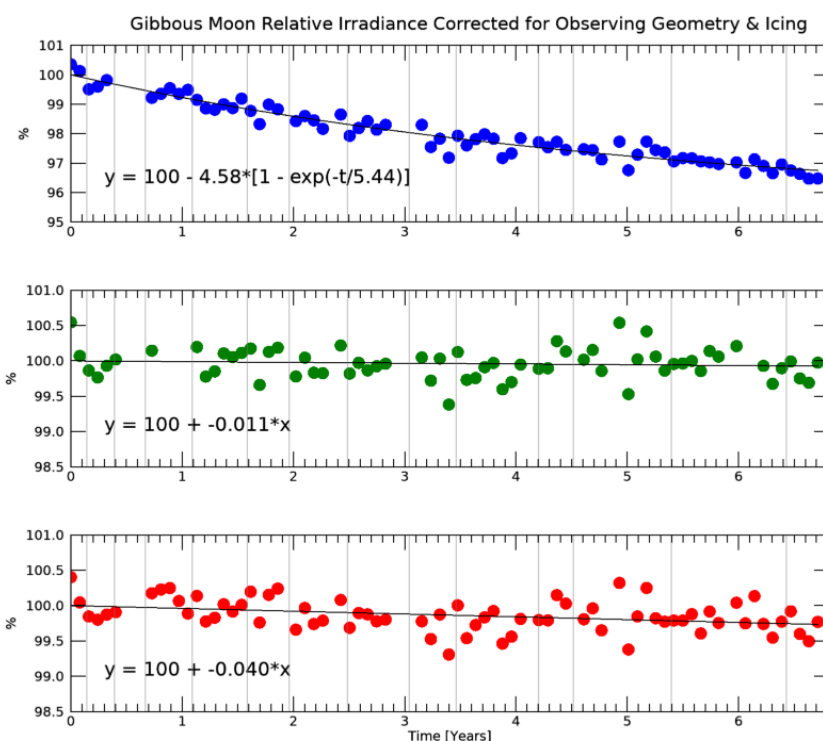


Figure 3-5. Lunar calibration time series after correction for observing geometry, including polarization effects, and icing in the ABO2 (top), WCO2 (middle) and SCO2 (bottom) channels indicate that the irreversible loss of radiometric throughput of the instrument is ~0.5%/year in the ABO2 channel (blue) and < 0.1%/year in the WCO2 (green) and SCO2 (red) channels.

3.5.3 OCO-2 Gain Degradation and Zero-Level Offset Correction Methodology

The change in responsivity of a sample or pixel at a given point in time is expressed in terms of a radiometric gain degradation factor, k :

$$k = \frac{1}{\text{Radiance}} * \sum_{i=1}^3 c_i (DN - DN_{dark})^i \quad (3-4)$$

where c_i are the preflight gain coefficients. Note that a k value less than 1 corresponds to a decrease in responsivity. The ARP contains coefficients that parameterize the k of each sample and pixel in the form of a polynomial function (degree ≤ 2) of time. These degradation coefficients are used together with the preflight gain coefficients in the conversion from the DN of L1a products to the radiance units of L1b products.

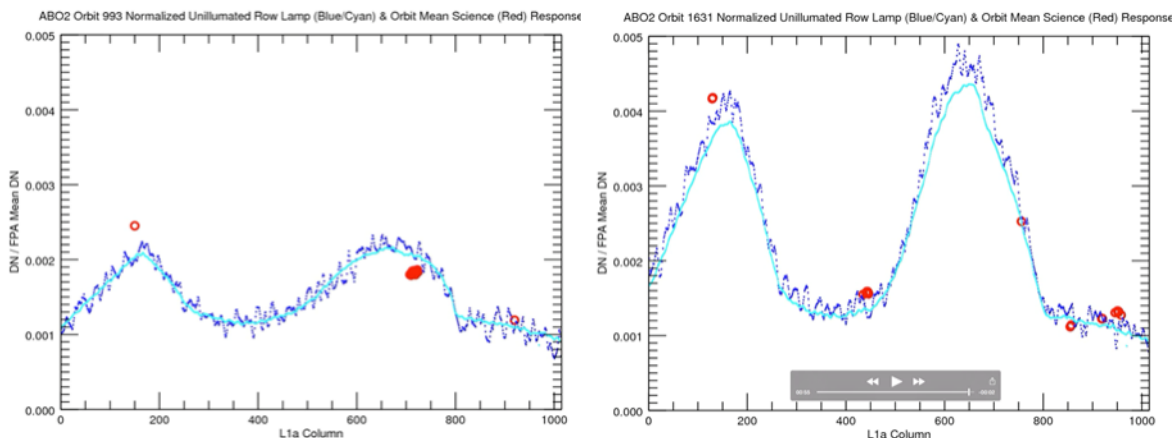


Figure 3-6. The zero-level offset, as measured in unilluminated rows in the shadow of the OCO-2 slit mask, early (left) and later (right) in a decon cycle. The curves represent smoothed and unsmoothed data from single pixel lamp calibrations, and the red dots represent color slice measurements acquired during normal science observations.

OCO-2 gain degradation factors are computed from the lamp, solar, and lunar calibration data sets. Both lamp and solar calibration sources use reflective, gold-coated diffusers, the reflectance of which can and does change over time. The lamps are the less radiometrically stable of the two over the long term, however, as they themselves undergo aging. For this reason, the solar calibration data set is used to determine the footprint mean radiometric gain degradation factor and spectral slope across the band. The solar calibration data set, on the other hand, is subject to higher order in-band variations due to solar Doppler shift. The lamp calibration data set is therefore used to determine the spectral shape of changes in responsivity within a footprint.

In a significant change from the v7 data set, v8 gain degradation factors were computed assuming that all change in solar calibration signal was due to calibrator degradation (Figure 3-7). The v7 data set overcorrected A-band radiances, thereby introducing an error in absolute radiance that grew with time. A study by the science team found that the retrieved albedo for desert regions increased at a rate of about 2.4% per year in v7 data, consistent with the error in the v7 calibration.

After the solar data have been corrected for the slow degradation, they are combined with lamp data to derive radiometric gain degradation factors that remove the effects of the icing, i.e., fast degradation. For the A-band, the gain degradation factors are spectrally, footprint, and time dependent (Figure 3-8). These gain degradation factors become progressively smaller as FPA icing reduces the responsivity of the spectrometer.

In v7, the gain degradation factors for the Weak and Strong CO₂ channels were set to unity. For v8, the Weak CO₂ and Strong CO₂ gain degradation factors were time dependent, and for v10 they are time and footprint dependent. There are small (sub 1%) residual spectral effects and

we will work to include them in the next data version. Another important change from v7 to v8 was to make the gain degradation time dependent rather than constant within each ARP range.

Figure 3-8 shows timeseries of radiometric gain for one spectral element for each channel.

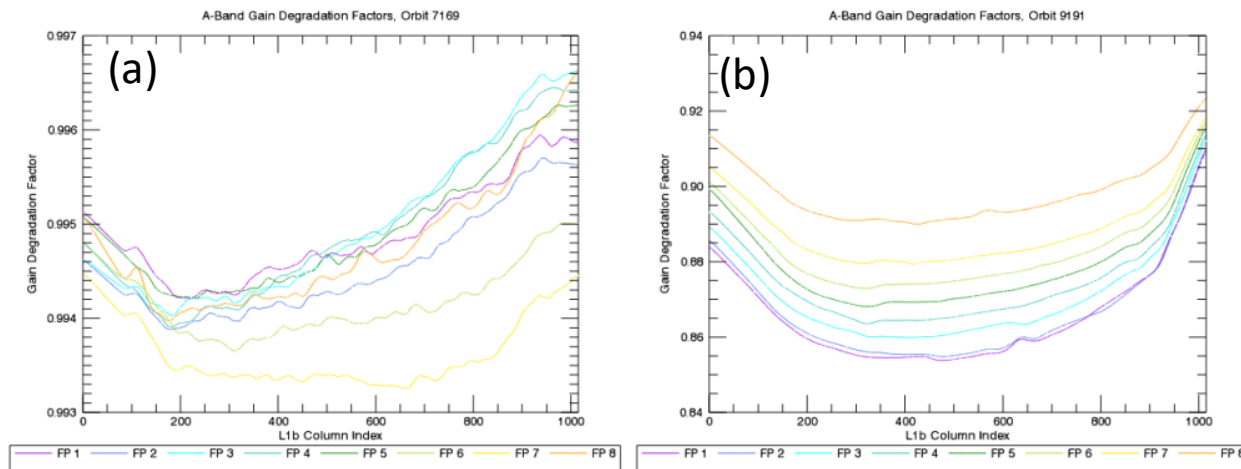


Figure 3-7. A-Band gain degradation factors from early (top) and late (bottom) in a fast degradation cycle.

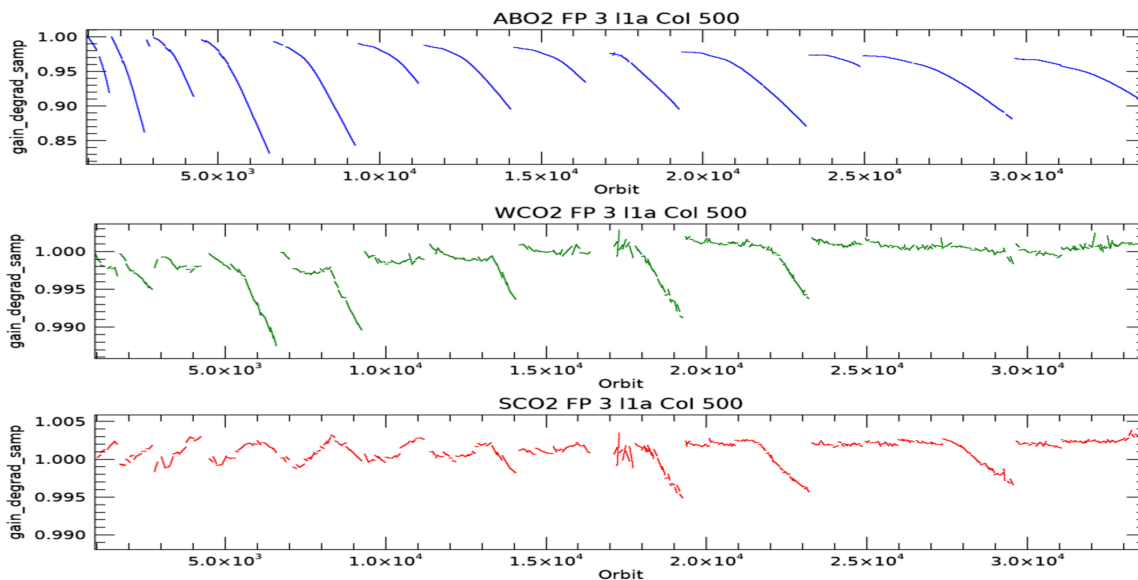


Figure 3-8. Time series of gain degradation factors for a single sample in each band. Note the range for the ABO2 is approximately 10x larger than for the WCO2 and SCO2 bands.

The overall impact of these changes is illustrated in Figure 3-9, where the radiance difference between v7 and v8 L1b are illustrated. Early in the mission changes are small for all bands. For the figure later in the mission, we see large ABO2 differences, whereas WCO2 and SCO2 changes are modest.

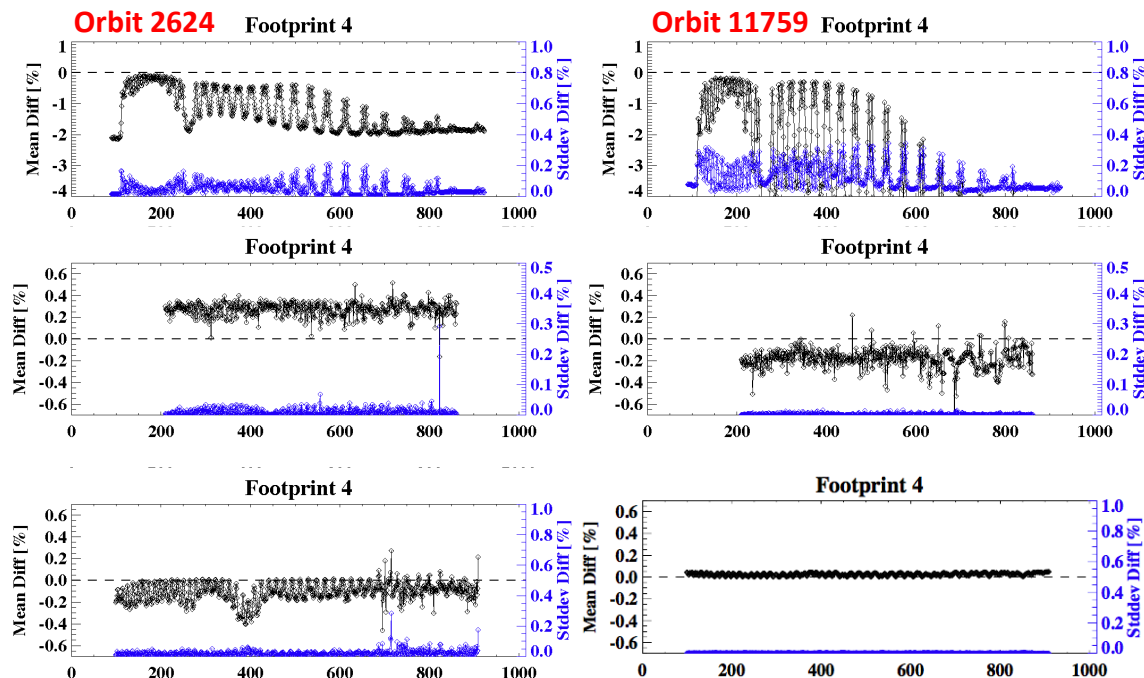


Figure 3-9. Differences between radiances from OCO-2 v8 and v7. The top panels are the band 1 radiance differences, middle panel are band 2 and bottom panels are band 3. The left panels are from an orbit early in the OCO-2 record (Orbit 2624, December 2014) and right panels are from September 2016 (Orbit 11759).

3.5.4 OCO-3 Gain Degradation

OCO-3 inflight radiometric calibration uses the same general approach as OCO-2, but derives separate gain degradation coefficients for all 1016×8 samples in all three bands. The most significant difference in methodology is that OCO-3 cannot use solar measurements to constrain changes in the primary lamp. For this reason, the secondary lamps are used more often, on 19% and 1% of calibration sequences, respectively.

All three lamps age with use, and as expected the lamp used most frequently has changed most since in-orbit checkout. Not all changes in lamp signal are attributable to lamp/calibrator aging, but additional sources are needed to separate this from true instrument degradation. For vEarly, all changes in Lamp 3 were attributed to gain degradation, leading to large errors when Lamp 3 signal decreased by roughly 15% in the first 300 days of the mission. For Build 10, we attribute changes in Lamp 1 to gain degradation, and changes in the other lamps to be lamp aging. This change is supported by Lamp 1 color temperature remaining constant, while Lamp 2 and Lamp 3 color temperature has decreased.

Lamp 1 is observed too infrequently to be used alone. We thus use an algorithm that mixes all three lamps. To have the best sampling in time, gain degradation k is the ratio of Lamp 3 radiance between flight and prelaunch divided by lamp 3 aging A_3 (Equation 3-5):

$$k(t_3) = \frac{L_3(t_3)}{L_3(pre) * A_3(t_3)} \quad (3-5)$$

Lamp 3 aging is determined from Lamp 2 aging instead of Lamp 1 directly to shorten the gaps in time for interpolation (Equation 3-6):

$$A_3(t_2) = \frac{L_3(t_2) * L_2(pre)}{L_3(pre) * L_2(t_2)} * A_2(t_2) \quad (3-6)$$

Lamp 2 aging is determined from Lamp 1, which is assumed not to age (Equation 3-7):

$$A_2(t_1) = \frac{L_2(t_1) * L_1(pre)}{L_2(pre) * L_1(t_1)} \quad (3-7)$$

As with OCO-2, the k values are fit as a function of time for every column, footprint, and band to determine gain_degrad_samp in the ARP.

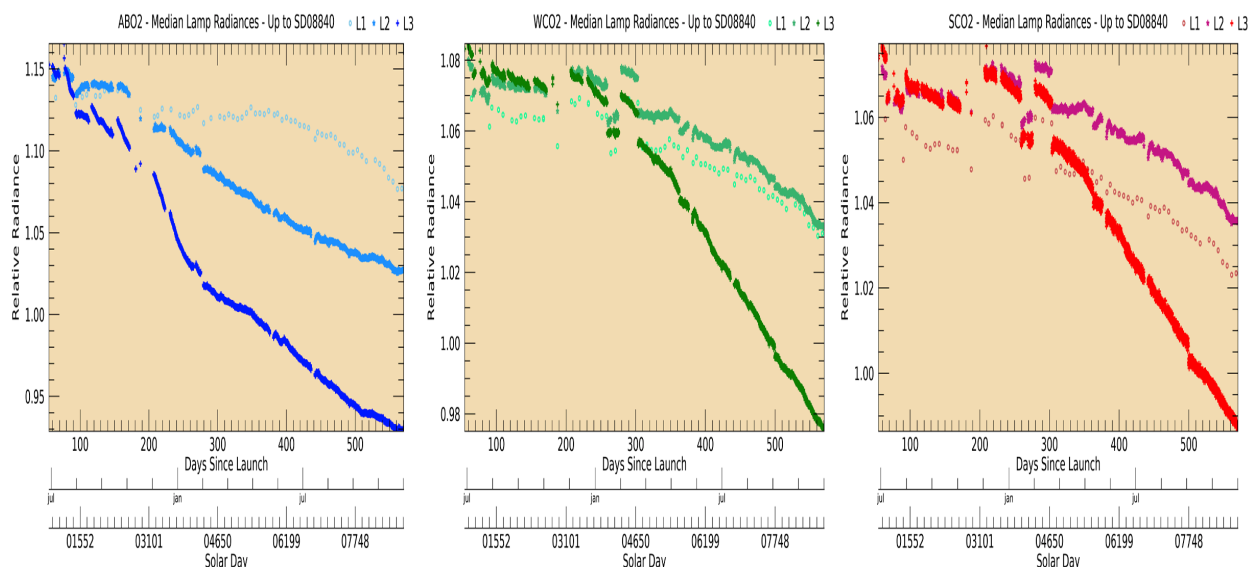


Figure 3-10. Trends of radiance relative to prelaunch for the OCO-3 ABO2, WCO2, and SCO2 channels with different colors indicating different lamps. Lamp 3, used most often, has aged most.

3.6 Zero-level Offset

In addition to variations in gain, the ABO2 (OCO-2 and OCO-3) and WCO2 (OCO-3 only) channels exhibit a small, spectrally dependent offset relative to the calibrated dark level. The amplitude of this zero-level offset increases the deeper into a fast degradation cycle the instrument gets and, for a given level of degradation, varies in proportion to the brightness of the scene. This offset appears to be due to light that is reflected off the FPA surface and then re-reflected and focused by the spectrometer optics back onto the FPA.

The existence of the zero-level offset was inferred from several lines of evidence, including the correlation of the time-dependence of the SIF correction factors, a subset of the L2 fitting parameters, and certain features in the L2 spectral residuals with the fast degradation cycles, and was confirmed by careful examination of the routine calibration data. Specifically, the zero-level offset was detected in the response of nominally unilluminated pixels in the shadow of the slit mask (Figure 3-6). As OCO-3 completed preflight testing after this was found in OCO-2 inflight

data, OCO-3 prelaunch radiometry includes corrections for this stray light, while OCO-2 does not.

Starting with OCO-2 v8, the ZLO correction is an A-band offset correction that is applied to the L1a data before it is converted to L1b (Figure 3-11). The ZLO correction rests on two assumptions. The first of these is that the ZLO signal is spatially uniform, not only within the science area but also across the nominally unilluminated rows that lie in the shadow of the slit mask. OCO-3 refined this assumption slightly by comparing unilluminated regions “above” and “below” the science area. The second is that the ZLO signal is proportional to the mean scene signal level.

The first assumption enables us to use the unilluminated pixel rows as proxies for the science area, where the small ZLO signal is difficult to disentangle from the much larger variations in radiometric response arising from the fast and slow degradation. Using these pixels introduces a complication of its own, however, in that being nominally unilluminated these pixels were never calibrated. As a result, the ZLO correction must be computed and applied before gain correction. Specifically, to compute the ZLO correction, the response of unilluminated pixels during routine single pixel lamp calibration is normalized by the mean pixel response over the illuminated portion of the FPA. This spectrally dependent single pixel ZLO correction is then converted into a spectrally and footprint dependent correction for application to summed data. To apply this correction to a given sample, the product of its time-dependent ZLO correction factor and the scene mean L1a signal is subtracted from its L1a response prior to conversion to L1b space.

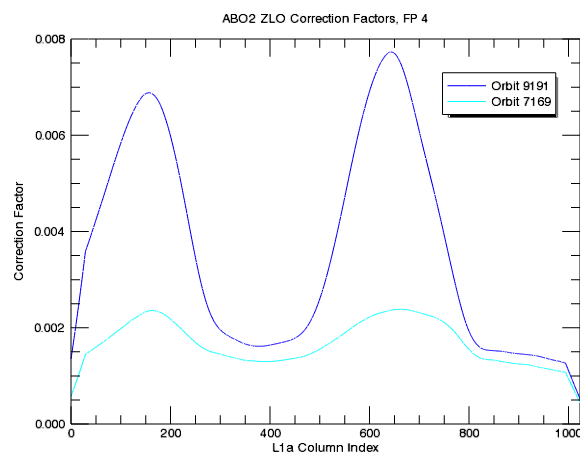


Figure 3-11. OCO-2 A-band ZLO correction factors near the beginning of (cyan) and deep into (blue) a fast degradation cycle.

Although these are small corrections ($< 0.9\%$ peak amplitude for the deepest fast degradation cycle), the correction did remove the signal from the L2 spectral residuals, flatten the SIF correction (for the 771 nm microwindow, not 757 nm), and remove decon related time dependence from other L2 variables.

OCO-3 ZLO is similar in character to OCO-2. The spectral shape is double-peaked, though the exact positions and relative intensity of the two peaks are slightly different. The OCO-3 ABO2 has unilluminated rows on both sides of the slit, while OCO-2 only has unilluminated

rows below footprint 1. OCO-3 also has unilluminated rows in the WCO2, below footprint 1 only, while the OCO-2 WCO2 band does not have any unilluminated rows.

3.7 Signal-to-Noise Ratio (SNR)

3.7.1 Calculation of SNR

The noise values are not stored directly in the L1B products, but they can be calculated using a few fields in the L1BSc data file and the following formulas for the noise equivalent radiance and consequently the SNR:

$$NEN = \frac{MaxMS}{100} \cdot \sqrt{\left| \frac{100 \cdot N}{MaxMS} \right| \cdot C_{photon}^2 + C_{background}^2} \quad (3-8)$$

$$SNR = \sqrt{\frac{100 N^2}{MaxMS * (C_{background}^2 \frac{MaxMS}{100} + C_{photon}^2 N)}} \quad (3-9)$$

where

- N —radiance value
- NEN —noise equivalent radiance (SNR is 1 at this radiance)
- $MaxMS$ —maximum measurable signal per band (see Table 3-5 below)
- C_{photon} —first coefficient of L1bSc/InstrumentHeader/snr_coef ([0, *, *, *])
- $C_{background}$ —second coefficient of L1bSc/InstrumentHeader/snr_coef ([1, *, *, *])

The user should be aware that the C_{photon} term captures the noise for a uniform scene – any additional variance that comes from a non-uniform scene causing changes in the instrument response based on pixel-to-pixel variability in either quantum efficiency or linearity is not included. Our experience has been that this approximation does not lead to significant errors in the Level 2 retrieval of XCO₂, but users investigating other questions should confirm that this approximation is valid for their application. Figure 3-12 shows an example set of coefficients.

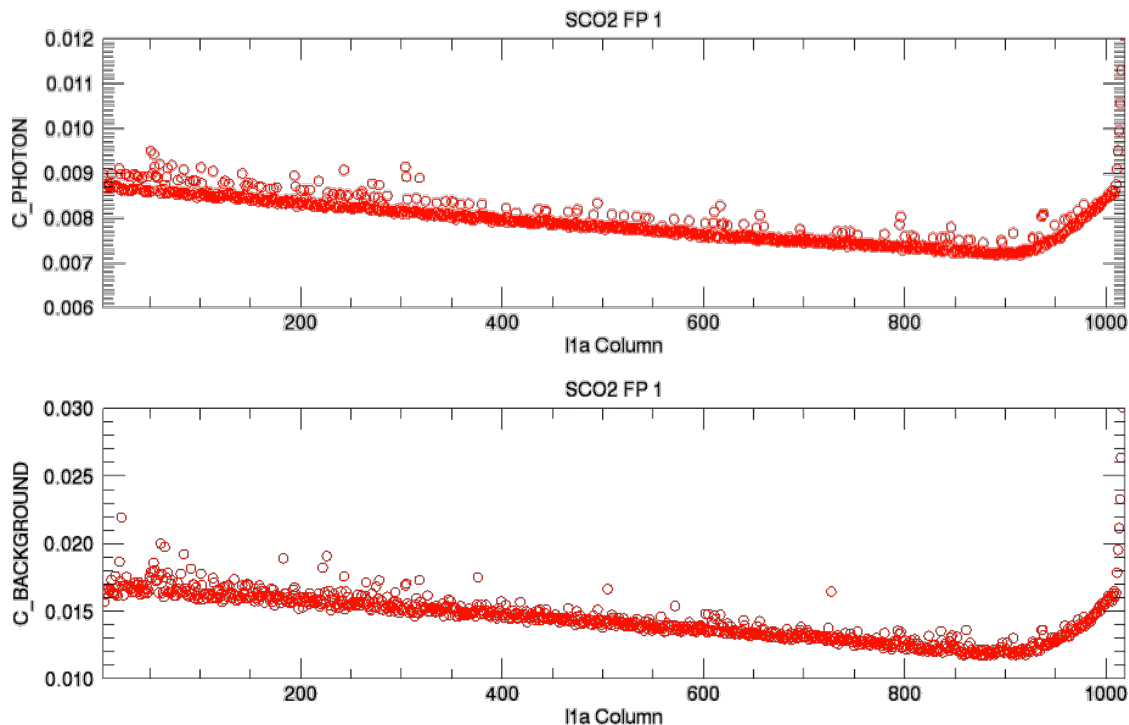
Table 3-5. OCO-2 and OCO-3 $MaxMS$ values.

Band	$MaxMS$ value (photons/m ² /sr/μm/s)
O ₂ A-band	7.00*10 ²⁰
Weak CO ₂	2.45*10 ²⁰
Strong CO ₂	1.25*10 ²⁰

For completeness, the “minimum measurable signal” values are included in Table 3-6. Note that these are requirements rather than measurements. The instruments measure signals as dim as MinMS with SNR values of at least 1 for 95% of spectral samples.

Table 3-6. Minimum measurable signal (*MinMS*) values.

Band	OCO-2 <i>MinMS</i> value (photons/m ² /sr/μm/s)	OCO-3 <i>MinMS</i> value (photons/m ² /sr/μm/s)
O ₂ A-band	7.50*10 ¹⁶	12.5*10 ¹⁶
Weak CO ₂	2.15*10 ¹⁶	3.50*10 ¹⁶
Strong CO ₂	2.15*10 ¹⁶	3.50*10 ¹⁶

**Figure 3-12.** Example of OCO-2 signal-to-noise coefficients.

3.7.2 Development of SNR Coefficients

For Versions 7–10, OCO-2 `snr_coef` were based on a combination of the data from the best five radiometric calibration tests during final instrument testing at JPL. While this model assumes a perfectly linear detector, which is not the case, it still fits the data reasonably well. The only significant errors observed are at extremely low signal levels, where this SNR expression overestimates the radiance uncertainty because we do not currently include a term for uncertainty due to dark subtraction but its presence is unlikely to significantly change the radiance uncertainty estimates.

3.7.3 In-Flight Updates of SNR Coefficients

In Version 11, the dark noise (background term) is now updated for each ARP range based on the dark calibration data. Preflight data continues to be used to derive photon noise, with one exception. We assume that the photon noise changes if the `gain_degrad_samp` term changes and adjust accordingly.

3.8 Bad Sample List

The Bad Sample List is stored directly in the L1bSc product in the /InstrumentHeader/bad_sample_list dataset. This dataset is used to identify bad samples that should be excluded by the retrieval algorithms. Entries in the array can range between 0 and 15 and is the sum of the following byte codes:

- 0 = Good
- 1 = Radiometric Problem (e.g., jumping pixel that needs to be identified for future bad pixel map, linearity problems)
- 2 = Spatial Problem (e.g., low wavelength 100 columns)
- 4 = Spectral Problem (e.g., ILS not trustworthy—last couple of samples on the CO₂ bands where the mask shadows the FPA)
- 8 = Polarization Problem (no occurrences yet)

3.9 Radiometric Footprint Bias and “Streak Flats”

In order to evaluate radiometric agreement among footprints, we selected scenes in the so-called “streak-flat” configuration for analysis. The OCO-2 instrument does not maintain its slit perpendicular to the track direction as it orbits the Earth, as shown on Figure 3-13. At every orbit, the slit slowly rotates and briefly reaches a configuration, which we term “streak flat”, where it is parallel to the ground track direction. At this point, the footprints are at their full length in the along track direction and very narrow cross-track. Also, some of them will nearly overlap, thus making these the best scenes to assess footprint biases.

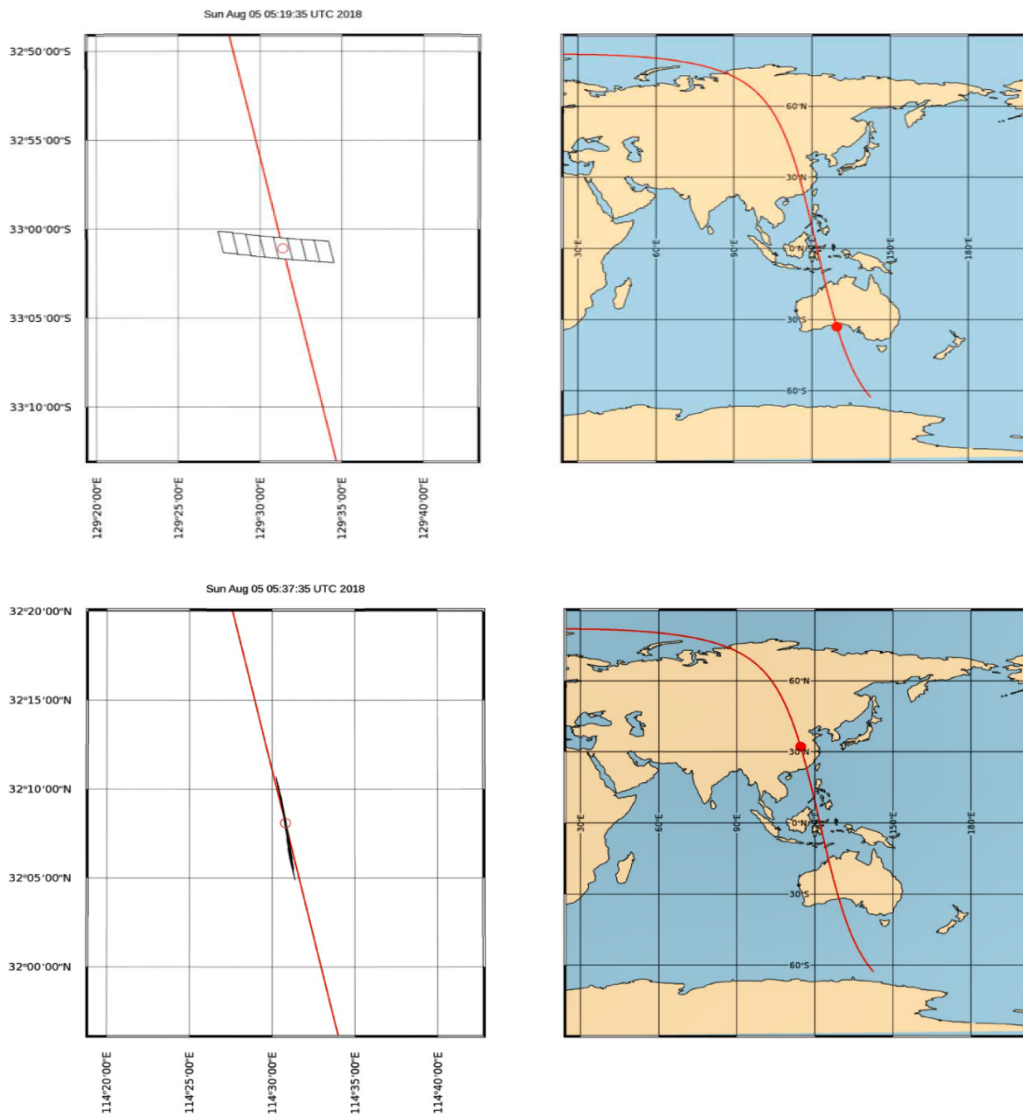


Figure 3-13. The two panels show the orientation of the slit with respect to the track direction at two different moments during an orbit. In the top panel, the slit is nearly perpendicular to the ground track direction and on the bottom panel, it approaches the streak flat configuration.

Figure 3-14 depicts how footprints 1, 3, 5, and 7 and footprints 2, 4, 6, and 8 of different frames around the streak flat frame (the one closest to the streak flat configuration) will nearly overlap. No overlap will occur between these two groups of footprints.

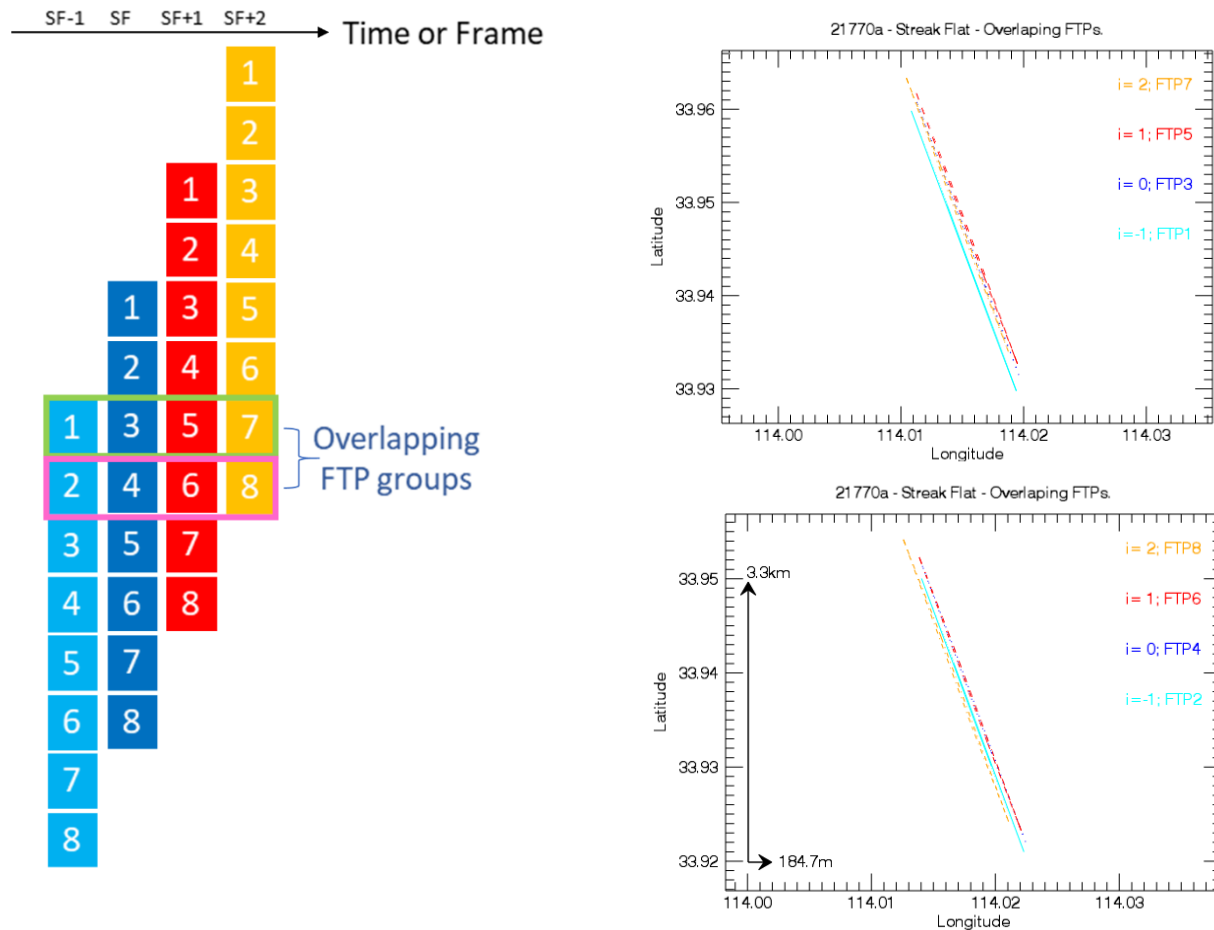


Figure 3-14. Footprints 1, 3, 5, and 7 and 2, 4, 6, and 8 from seven frames around the streak flat frame nearly overlap. On the left panel, representations of the slit, with its 8 footprints, are shown for consecutive frames. Each frame is shown as a different color and the vertical relative positions represent the physical position along track. The left panels shown the nearly overlapping footprints mapped on the ground. In the legends, i refers to the frame distance to the streak flat frame ($i=0$).

For this analysis, we selected seven frames per granule, centered on the streak flat frame, which allowed four intercomparisons among footprints for each of the two footprint groups. We used 5815 L1bSc granules, selected from QTSv10 orbits. In each granule, the streak flat frame was selected to be the one with the smallest angle between slit and ground track and also required to be less than 0.03 degrees.

We defined continuum regions (as shown on Figure 3-15, ABO2: 0.7707–0.7711, WCO2: 1.60567–1.60593, SCO2: 2.06015–2.0603 μm) where the median radiances were computed and compared among footprints. The long wavelength edges of the Weak and Strong bands were not used due to large amounts of bad samples in those regions, which vary with footprint.




Figure 3-15. Selected continuum regions: ABO2: 0.7707–0.7711, WCO2: 1.60567–1.60593, SCO2: 2.06015–2.0603 microns. Footprint 7, frame 500 of granule 21342a is shown.

Results were produced for three separate periods: orbits 958–7260, 7261–17785, and 17785–31462. For each footprint of each band, for each of these three periods, a factor was derived based on the mean value of gaussian fits to radiance/(group radiance) histograms. These are shown on Figure 3-16 as colored circles. One standard deviation is also shown in color, along with the normalized histograms in grayscale.

Unfortunately, while these factors were incorporated into gain_preflight_samp for B11, an implementation error in gain_degrad_samp led to the changes being entirely counteracted. Therefore, Versions 10 and 11 have identical footprint biases. This does not impact the validity of the analysis in this section.

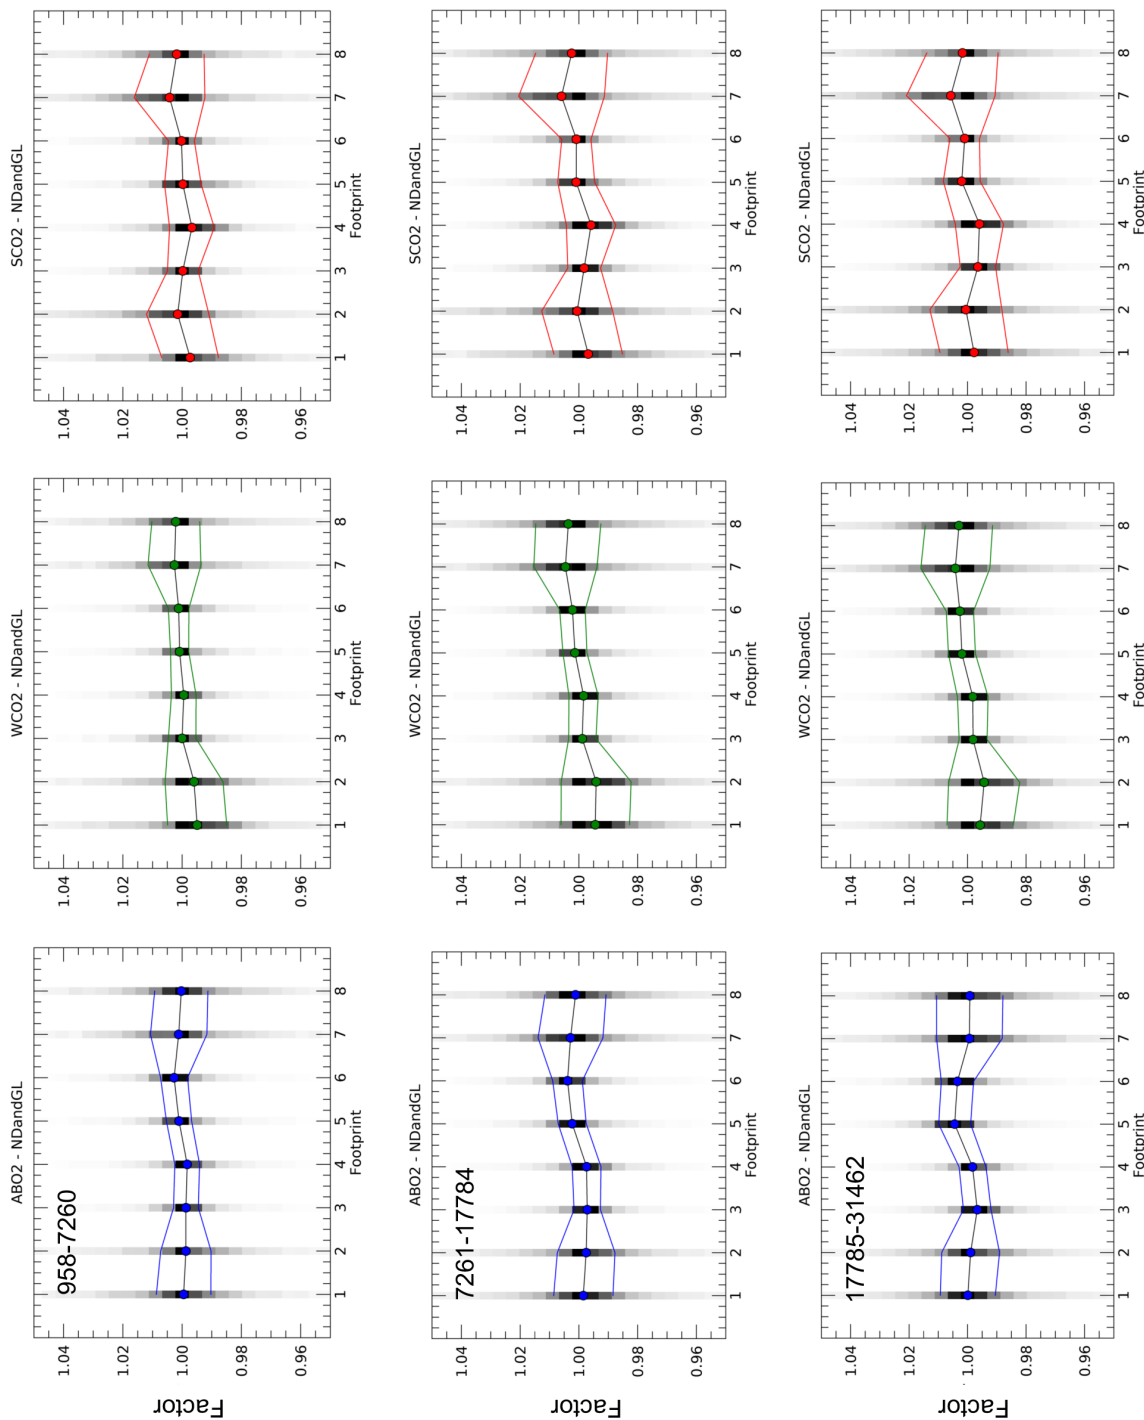


Figure 3-16. Factors (colored circles) derived from the streak-flat analyses for each footprint, of each band, for three time periods, used to divide pre-launch gains. In gray scale, normalized histograms of radiance/(group radiance) are shown. Colored lines indicate one standard deviation.

4 Spectral Calibration

The spectral dispersion and ILS of the OCO-2 instrument were well characterized and calibrated prior to launch (Frankenberg et al. 2015; Lee et al. 2017). The dispersion relation for each spectrometer specifies the mapping of the wavelengths of light onto specific pixels on the FPAs. Small changes in the dispersion coefficients were expected after launch due to Doppler shifts of the observed radiation (i.e., relative motion between the spacecraft and reflecting target). Dispersion changes can also be introduced by small displacements of the (cryogenically isolated) FPAs relative to the OBA, or small changes in the tilt of the gratings introduced by thermal gradients across the OBA. Dispersion coefficients are provided in the L1bSc files (dispersion_coef_samp). In the L2 retrieval of geophysical quantities, the Doppler shift is accounted for and the dispersion is corrected for each sounding.

The ILS describes the spatial distribution of the light on the FPA that is produced when light from the slit is collimated, dispersed into a spectrum by the grating, and reimaged on the FPA. In the OCO-2 instrument, the ILS changes from footprint to footprint, and across each summed spectrum. The OCO-2 ILS was characterized during pre-launch testing, using techniques similar to those used for the original OCO instrument (Day et al. 2011; Frankenberg et al. 2015; Lee et al. 2017). In principal, changes in the focus of spectrometer collimator or camera assemblies could also produce changes in both the dispersion and the ILS, but no ILS changes have been verified on orbit. The ILS information is contained in the L1bSc files in two key fields: ils_delta_lambda and ils_relative_response. ILS are covered in more detail below, and their use is described in the Data User's Guide.

4.1 Spectral Dispersion Coefficients

The dispersion coefficients express the relationship between the spectral element index (an individual FPA column or spectral sample) and its associated wavelength (see Figure 4-1). These data are contained in L1bSc/InstrumentHeader/dispersion_coef_samp. The dispersion equation is the same for OCO-2 and OCO-3. The coefficients are used as follows:

$$\lambda = \sum_{i=0}^5 c_i \cdot \text{column}^i \quad (4-1)$$

Here, column refers to the column number in the L1bSc files (1 to 1016), which is raised to the power i .

An example calculation of the wavelength grid is:

$$\begin{aligned} \lambda = & 0.757633 + 1.75265 \times 10^{-5} \cdot \text{column}^1 \\ & - 2.91788 \times 10^{-9} \cdot \text{column}^2 + 3.29430 \times 10^{-13} \cdot \text{column}^3 \\ & - 2.72386 \times 10^{-16} \cdot \text{column}^4 + 7.66707 \times 10^{-20} \cdot \text{column}^5 \end{aligned} \quad (4-2)$$

The spectral dispersion and ILS were characterized and calibrated prior to launch (Frankenberg et al. 2014; Lee et al. 2017). In flight, there are small changes in the dispersion due to Doppler shifts and thermo-mechanical effects associated with the cryo-cooler drive level, and OBA thermal gradients, which change between Decontamination cycles (Crisp et al. 2017).

These small dispersion changes can be corrected by adjustments in the wavelength “offset” and “stretch” coefficients in the dispersion polynomial (terms 1 and 2 in the 5th order polynomial), that are derived from the spectral fitting process in the L2 retrieval algorithm. In the OCO-2 V8 product the retrospective L1B dispersion is modified to include these new coefficients.

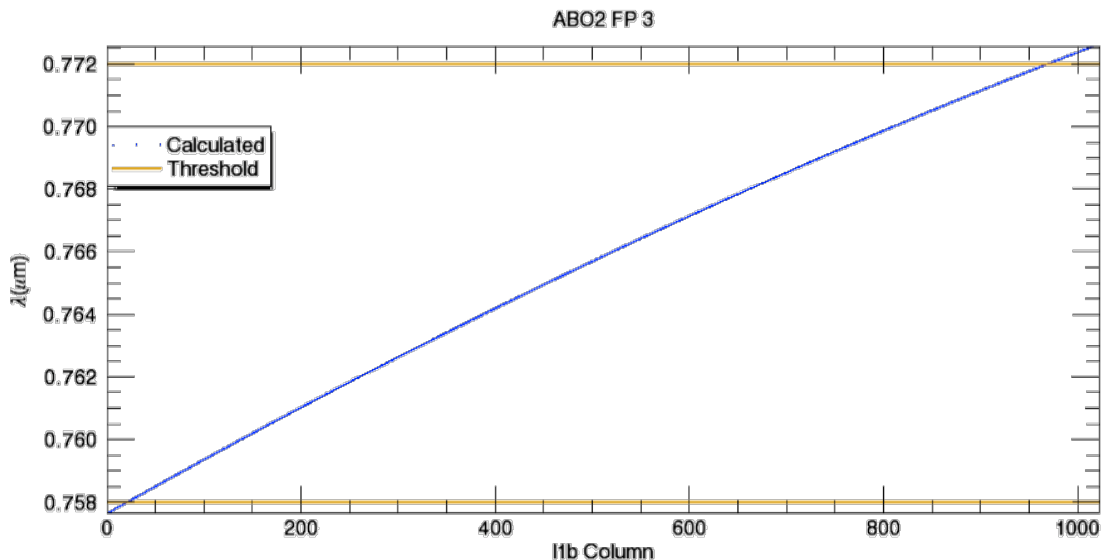


Figure 4-1. An example of the conversion of the focal plane column to wavelength scale. The blue line shows the column number-to-wavelength transformation. The yellow lines indicate the required wavelength range for this band.

4.2 Instrument Line Shape

The central challenge of the OCO-2 spectral calibration is determining not just a single ILS, but rather the ILS for every spectral pixel index, footprint, and band. The three bands, eight footprints per band, and 1016 spectral samples per footprint yield 24,384 individual ILS functions. This is in contrast to, e.g., the Thermal and Near Infrared Sensor for Carbon Observations-Fourier transform spectrometer (FTS) aboard the Greenhouse Gases Observing Satellite (GOSAT), which in theory has just two ILS functions (one per polarization), as described by Yoshida et al. (2012). However, the physics of the OCO-2 instrument design dictates that the ILS and centroid wavelength response (dispersion) of OCO-2 should vary smoothly in the spectral dimension across each band. The details of the measurements and analysis for determining the ILS are reported in Lee et al. (2017).

In the data product, for each band, footprint, and spectral element ($3 \times 8 \times 1016$ samples), there are two 200-element lookup tables: `L1bSc/InstrumentHeader/ils_delta_lambda` and `L1BSc/InstrumentHeader/ils_relative_response`. These curves describe the response of each spectral element of the instrument versus wavelength, and can be used to convolve high spectral resolution spectra for comparison with OCO-2 spectra.

Initial determination of the ILS was performed using tunable diode lasers that were stepped through a range of wavelengths covering the OCO-2 spectral range. The final ILS was then optimized and validated by comparing solar spectra recorded simultaneously by OCO-2 and an FTS, as described by O’Dell et al. (2011). As reported in Lee et al. (2017), the resulting ILS

profiles (Figure 4-2) showed agreement between the two spectra to approximately 0.2% rms, satisfying the preflight calibration requirement of < 0.25% rms.

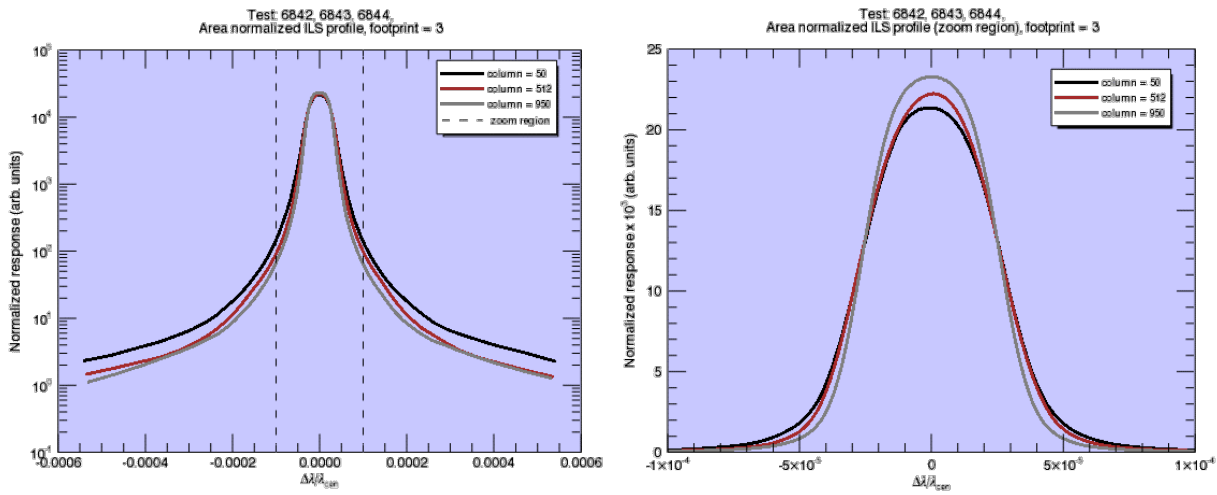


Figure 4-2. Example OCO-2 ILS profiles for the O₂ A band at three different spectral pixel indexes. (Left) Semi-log plot of the ILS. (Right) Linear plot of the ILS core, demarcated by the dashed vertical lines in the left plot.

Post launch OCO-2 spectral residuals in the retrievals have been carefully analyzed for evidence of ILS errors, using on-orbit data. The solar calibration observations have provided an opportunity to identify and trend relative changes in the ILS. However, since the flight diffuser may make subtle changes to the ILS, an absolute measurement may not be possible. To date, these on-orbit measurements provide no compelling evidence of changes in the widths or shapes of the ILS functions derived prior to launch.

In B11 a correction was made to the TVac radiometry to account for the Zero Level Offset (ZLO, described in Section 3.6), which produced a moderate improvement in the ABO2 spectral calibration by lowering the wings of the ILS. ZLO was first discovered on OCO-2 in the on-orbit data. It is caused by a small amount of light, initially reflected off a thin layer of ice on the focal planes, being re-reflected and focused by the optics back onto the detectors. This causes a small offset to the radiometry, and is most apparent in the ABO2 since its anti-reflection coating has an index of refraction most similar to that of ice. No significant improvement was observed in the WCO2 and SCO2 bands, possibly due to the reduced influence of ice on these detectors, and the fact that their spectra exhibit shallower absorption lines, less affected by the ZLO correction.

For OCO-3, similar TVac measurements were made in April 2018 and an initial ILS and dispersion were delivered in 2019 for use in vEarly. The spectral calibration of OCO-3 was slightly worse than OCO-2, mainly in the SCO2 due to lower SNR in the laser-based measurements, nevertheless it was deemed sufficient for on-orbit use. As with OCO-2, a ZLO correction will be made to the TVac radiometry with anticipated, similar improvement to ABO2 and WCO2 residuals. The new ILS will be included in B11, and results will be reported in a forthcoming manuscript (Lee et al. 2023).

5 Geometric Calibration

To retrieve XCO₂ estimates with accuracies near 0.25%, the length of the atmospheric optical path traversed by the sunlight must be known to a comparable accuracy. An accurate geolocation of each sounding footprint and an accurate estimate of the mean elevation of the footprint are essential for defining the optical path length associated with each spectrum. In addition, because the OCO-2 and OCO-3 spectrometers are imaging spectrometers, the measured radiances across each spectral range must be referenced to the same spatial coordinates on the target. Any rotation of the FPA columns relative to the grooves on the diffraction grating or the spectrometer slits (“clocking”) or optical artifacts such as smile or keystone can compromise the spectral-spatial registration. The following two subsections summarize the geolocation, and spatial-spectral calibration approaches used for OCO-2 and OCO-3.

5.1 Pointing and Geolocation

To establish the geolocation of each sounding footprint, the location of the spacecraft along the orbit track, the pointing of the instrument boresight relative to a local coordinate system, the relative pointing of the fields of view of the eight footprints in the three spectrometer channels and the mean topographic elevation of the footprint must be known. For OCO-2, where the spacecraft points the instrument, the key challenge is to establish the pointing of the instrument boresight relative to that of the orbital coordinate system. Because the boresight of the OCO-3 instrument is determined by its 2-axis pointing mirror assembly (PMA), the relative pointing of the PMA and spacecraft coordinate system must be known as well.

For OCO-2, the location of the observatory along its orbit track is determined by a Global Positioning System (GPS) sensor. The spacecraft coordinate system was adopted as the local coordinate system, and its orientation relative to the fixed stars is determined from star tracker measurements. The relative pointing of the star tracker and the instrument boresight were measured prior to launch, and these measurements were refined during after launch using observations of the lunar disk. The spatial fields of view of each footprint and the relative boresight alignment of the three spectrometer slits were characterized during the pre-launch instrument tests and refined once in orbit using observations of the lunar disk. The geolocation algorithm combines this pointing information with a spacecraft ephemeris and attitude derived from the GPS and the star tracker, and an accurate, high-resolution digital elevation model (DEM) to estimate the intersection between the pointing vector for each footprint and the Earth’s surface. The end-to-end performance of the pointing and geolocation algorithms were verified using observations of distinct surface features, such as coastline crossings (Crisp et al. 2017).

The pre-flight requirement for pointing accuracy was one milliradian (0.057°). Compliance with this requirement was demonstrated during the in-orbit check out period. However, analysis of the V8 data products revealed significant variations in XCO₂ in regions with topographic changes (Wunch et al. 2017; Kiel et al. 2019). For example, XCO₂ estimates retrieved from target observations over the Lauder TCCON site showed biases as large as 3 ppm, which were spatially correlated with the 150–200 m hills just south of the TCCON station (Figure 5-1). Similar variations were seen in routine nadir and glint observations over regions with significant elevation variations, such as Death Valley, California, and the Atacama Desert, Chile (Kiel et al. 2019). An investigation traced these biases to a small (0.03°) pitch error in the pointing of the instrument boresight. With this information, the DA team developed a bias correction that could

be applied to the V8 XCO₂ estimates after the fact, precluding the need for a full reprocessing campaign. This bias correction reduced the slope-dependent XCO₂ biases to <1 ppm in the V9 “lite file” products, which were delivered to the DISC starting in October 2018.

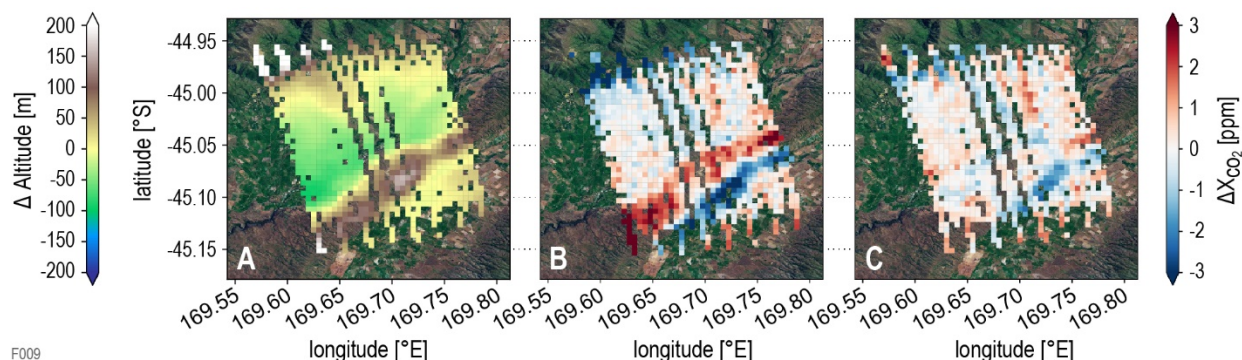


Figure 5-1. (A) topographic elevation variations near the Lauder TCCON station. (B) Estimates of V8 XCO₂ retrieved from OCO-2 observations collected near the Lauder station and in other regions with significant topography showed anomalous variations that were strongly correlated with topographic slopes. (C) Correcting the pointing error in the V9 product reduces XCO₂ biases by more than half, illustrating the impact on XCO₂ products of continual improvements to L1b products.

5.2 Radiance Discontinuities Due to FPA Rotation—The Clocking Correction

For imaging spectrometers like those used by OCO-2 and OCO-3, the spectral radiances across the entire spectral range of all three spectral channels must originate from the same spatial footprint within the instrument’s field of view. This objective is most easily achieved by an imaging grating spectrometer if the spectrometer slits, the grooves on the diffraction gratings, and columns of the FPAs are parallel to each other, and the spectrometer optics introduce negligible optical aberration (i.e., smile, keystone) between the slits and the FPAs. A given row on the FPA will then correspond to a given spatial location within the field of view.

For the OCO-2 spectrometer channels, the slits and gratings are well aligned and there is very little optical aberration between the slits and FPA. However, perfect alignment of the FPAs was not possible due to a physical obstruction discovered late in the instrument assembly process. The focal plane arrays are therefore slightly rotated (or “clocked”) with respect to the slit and grating (Figure 5-2). Consequently, a given geographic position does not map onto a single row of pixels across the entire spectral range sampled by the FPA. Instead, it changes from row to row (roughly linearly) with spectral position (i.e., column). This is not an issue in spatially homogeneous regions but can introduce spectral radiance discontinuities for scenes with substantial variations in illumination across a given spatial sample.

Identifying and correcting these discontinuities is complicated, somewhat, by the fact that the instrument sums ~20 spatial pixels into each spectral sample that is transmitted to the ground. For example, in lunar calibration data, where radiances can decrease from a maximum value to the noise level over a span of four or five pixels and where there are no obscuring absorption features, discontinuities are clearly discernable above the noise level in nearly every measurement. Methods for identifying and correcting these discontinuities are described in the following section. For OCO-3, only the A-band will require declocking.

Because fields of view must ultimately be mapped onto discrete pixels, the start and end positions of each footprint are adjusted in single pixel increments (i.e., clocking jumps) in row space at selected columns (Figure 5-2). These jumps can introduce discontinuities in the mean measured radiances for a spatial footprint if there are sharp gradients in the illumination near its boundaries. To compensate for these discontinuities, we have developed a simple clocking correction algorithm that enforces continuity in radiance across single pixel footprint shifts. The principal idea behind the algorithm is that although we do not have access to the full matrix of radiances over all rows and columns, we can use the entire row (i.e., spatial) information of a subset of the twenty color slices as a proxy for the spatial variation of intermediate columns, particularly those columns where clocking jumps occur. More specifically, at a small number of color slice positions we compute what the magnitude of the discontinuity in radiance would be if the slice position coincided with a particular clocking jump. We then make the assumption that this discontinuity is a valid approximation over some spectral neighborhood of the color slice and that the results from all the color slices can then be smoothly joined in some way.

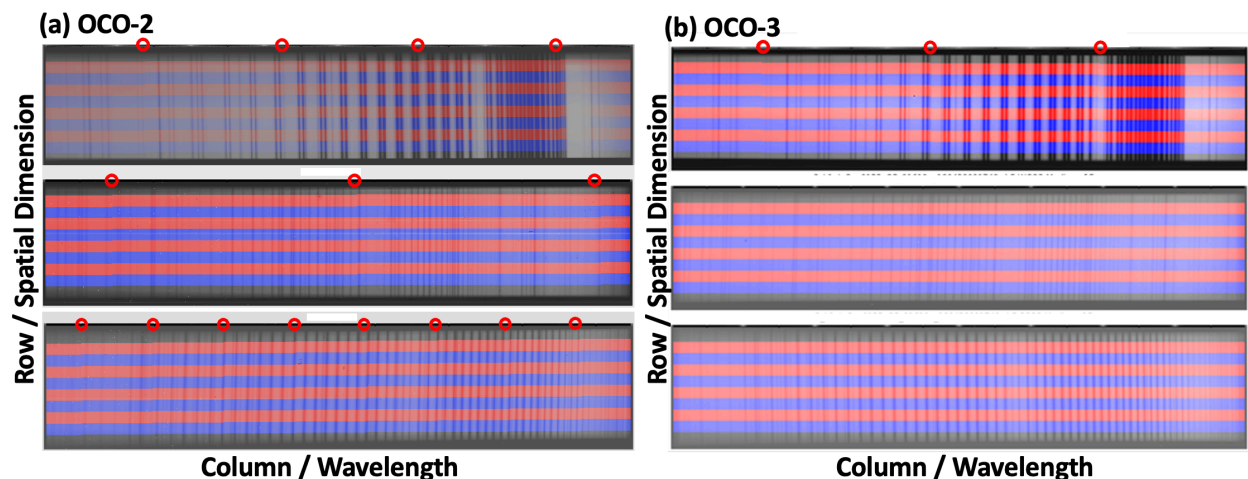


Figure 5-2. (a) Images of atmospheric spectral from the OCO-2 ABO2 (top), WCO2 (middle) and SCO2 (bottom) channels showing the rotation, or “clocking” of the FPA columns with respect to the dark O_2 and CO_2 absorption lines. The 8 spatially-summed footprints in each channel are shown from bottom (footprint 1) to top (footprint 8) in blue and red shading. The clocking is most obvious in the ABO2 channel, where the O_2 lines are tilted slightly counterclockwise and in SCO2 channels, where the CO_2 lines are tilted clockwise. The columns where the clocking corrections are applied in each channel are highlighted with red circles. (b) Same as panel (a) for the OCO-3 instrument. For OCO-3, only the ABO2 channel needs a clocking correction because the WCO2 and SCO2 channels are well aligned.

The selection of which columns comprise the set of color slices can be changed by ground commands. One could place color slices adjacent to the columns with the clocking jumps. However, this is problematic for this application because some of these columns include sharp atmospheric or solar absorption features, and these features can be shifted by 2–3 columns by Doppler shifts associated with the spacecraft motion. If a color slice is placed in a region where the radiance changes rapidly (e.g., near an absorption line), the variation of radiance along a given row may be more indicative of the local spectral structure than of the spatial features of the measurement. Therefore, it is very important to place color slices where spectral variation is limited, i.e., in continuum regions. The availability of pixels in clear continuum regions varies

substantially from band to band. For the O₂ A-band, continuum regions were selected just beyond the measurable high J value lines of the P and R branches as well as between the two branches. A similar approach was adopted for the weak CO₂ band near 1.61 μm . In contrast, the deeply saturated strong CO₂ band near 2.06 μm offers no spectral regions as transparent as in the other two bands, and none of the color slices used by the clocking algorithm lies in between the P and R branches.

Figure 5-3 through Figure 5-9 illustrate the nature of the clocking error and the choices of color slices for each of the three OCO-2 bands. For each band, there is a large figure which contains three different graphic elements. First are the nearly horizontal solid lines which delineate the boundaries of the eight footprints. The discrete jumps in the footprint boundaries illustrate the onboard software clocking. Note that the magnitude of the clocking effect varies among the bands, ranging from 8 discontinuities in the most severe case (SCO2) to 2 in the least (WCO2). Second, the dashed vertical lines indicate the current selection of all twenty color slices (as of October 2014). *Only those slices that are labeled by an index are used by the clocking algorithm.* Finally, a representative spectrum is superimposed to provide some context for the color slice selection.

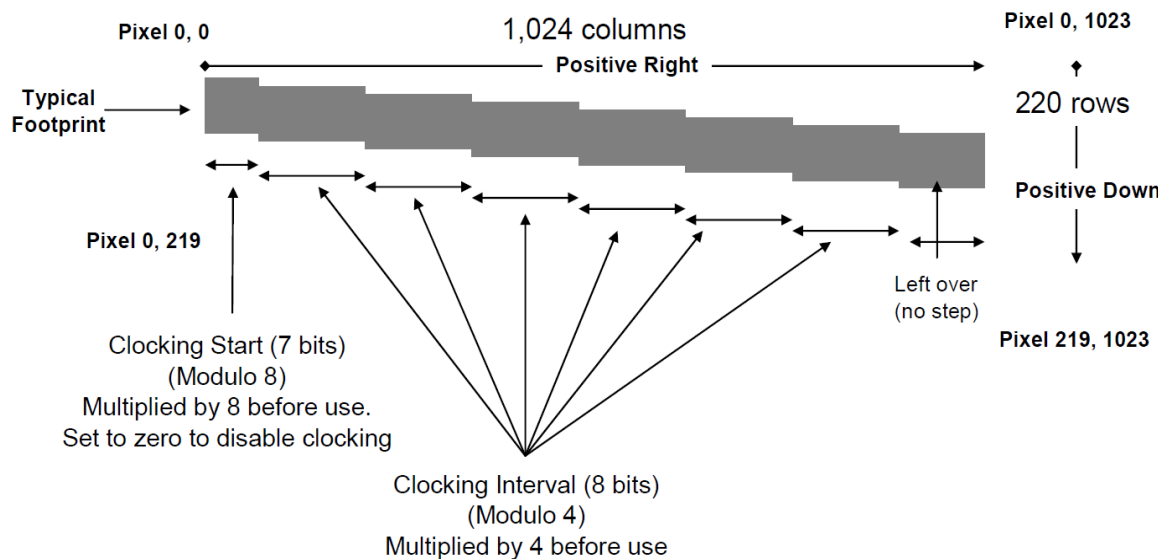


Figure 5-3. The onboard software "clocking" scheme used to sample the same spatial information (vertical dimension) across the entire spectral range (horizontal) for a rotated FPA.

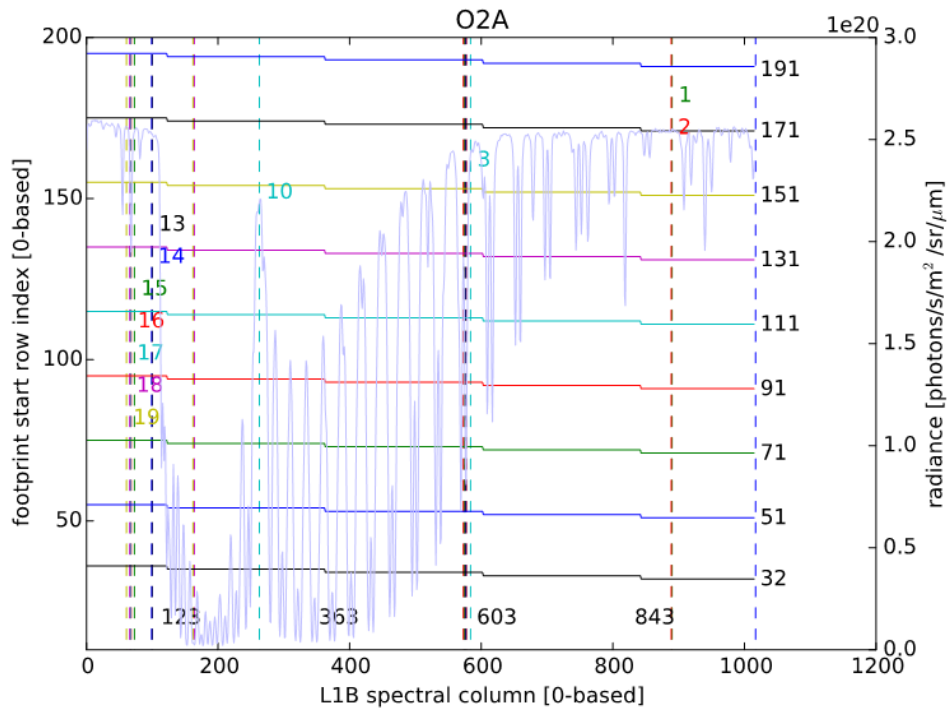


Figure 5-4. Schematic illustrating the clocking for the OCO-2 O₂ A band along with the current color slice selection. Horizontal, piecewise-constant solid curves delineate footprint boundaries. Dashed vertical lines indicate color slice selection. Only those that are numbered are used for clocking. Finally, a typical spectrum of the band is overlaid for reference.

Generally, the color slices used by the clocking algorithm are located in spectral regions with little absorption. Not all twenty color slices are used for clocking; other color slices have been located within strong absorption features to assess spatial inhomogeneity due to clouds and aerosols or to monitor solar lines. In the O₂ A-band, slices 16, 17, and 18 lie in the neighborhood of a solar line and were mistakenly included in the clocking correction. However, it was found that corrections using these slices correlate very strongly with those using slices 15 and 19, which do lie in the continuum, so that the inclusion of these color slices in the clocking correction algorithm is unlikely to have significant impact. Slices 3, 4, 5, and 15 from the SCO2 band may in the future be eliminated or replaced, since they lie close to weak absorption lines. Further testing is required to determine what difference their omission would make.

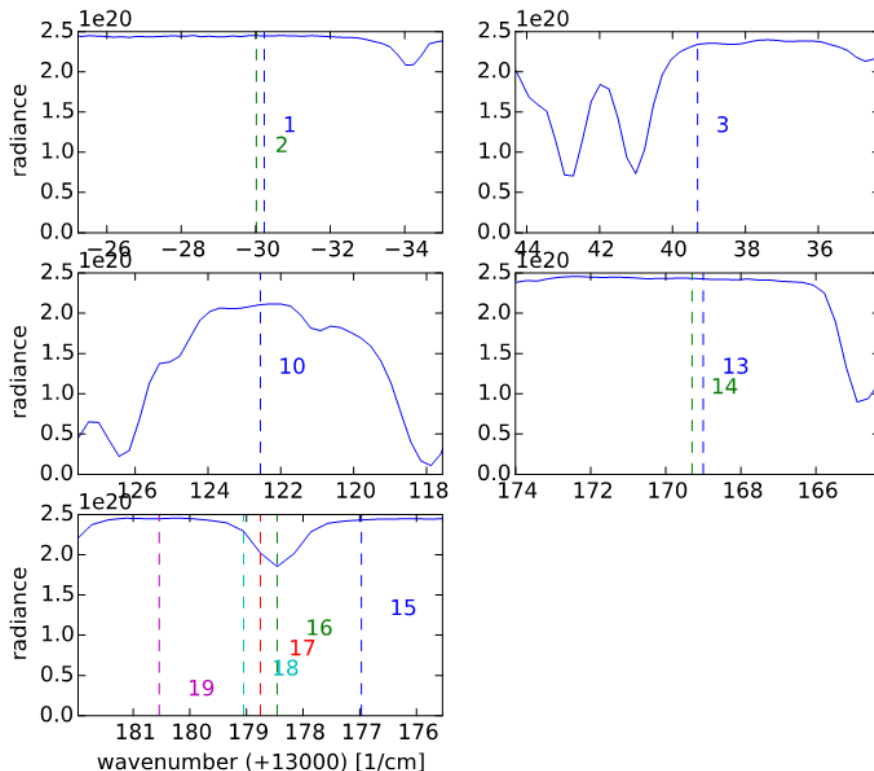


Figure 5-5. The color slices used in the clocking correction for the O₂ A band (shown in Figure 5-4) are aggregated into five “groups.” Each panel depicts a single group.

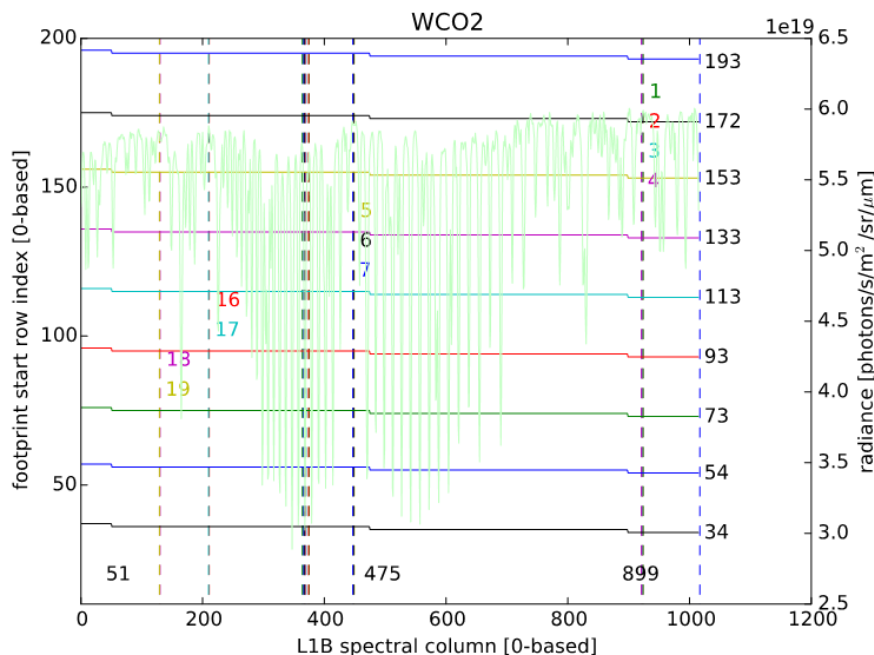


Figure 5-6. Schematic illustrating the clocking for the WCO2 band along with the current color slice selection. Piecewise-constant solid curves delineate footprint boundaries. Dashed vertical lines indicate color slice selection. Only those that are numbered are used for clocking. Finally, a typical spectrum of the band is overlaid for reference.




Figure 5-7. The color slices used in the clocking correction for the WCO2 band (shown in Figure 5-6) are aggregated into four “groups.” Each panel depicts a single group.

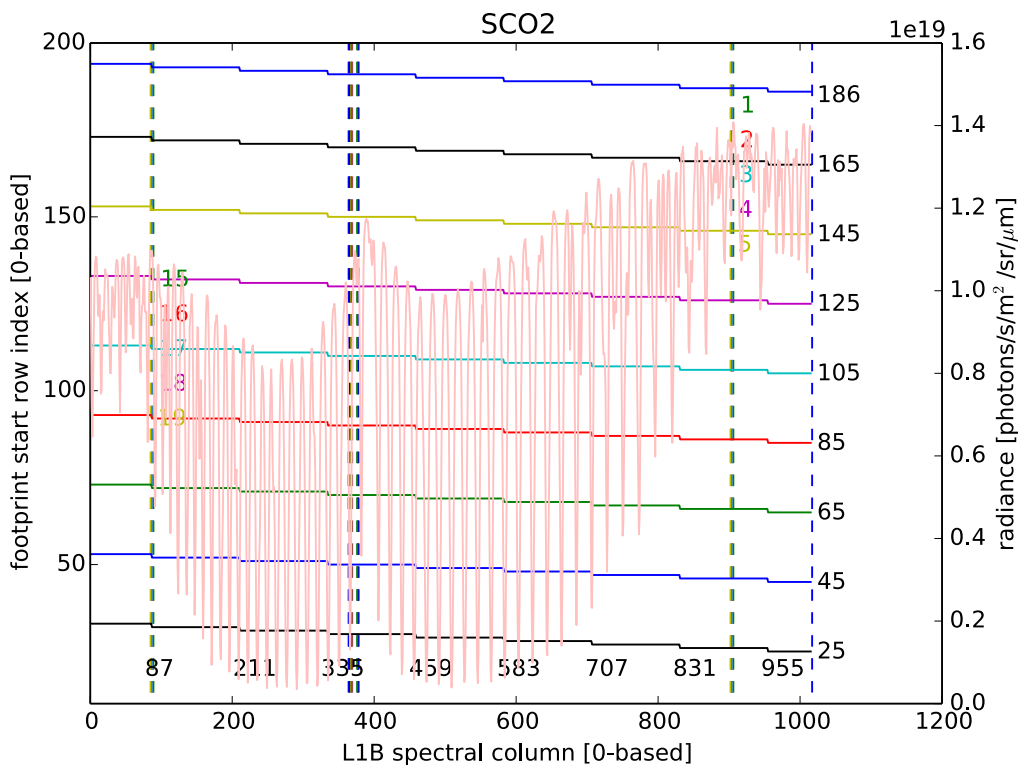


Figure 5-8. Schematic illustrating the clocking for the SCO2 band along with the current color slice selection. Piecewise-constant solid curves delineate footprint boundaries. Dashed vertical lines indicate color slice selection. Only those that are numbered are used for clocking. Finally, a typical spectrum of the band is overlaid for reference.




Figure 5-9. The color slices used in the clocking algorithm for the SCO2 band (shown in Figure 5-8) are aggregated into two “groups.” Each panel depicts a single group.

5.2.1 Clocking Algorithm Description

5.2.1.1 First Stage: Calculating Average Radiances

The algorithm used to compute modified L1B radiances proceeds in two stages. In the first stage, for each footprint-color slice pair, (f, c) , average radiances over all possible footprint start and end positions are computed and saved to a file:

$$\bar{R}(f, c, j) = \frac{1}{M_{f,c}} \sum_{r \in FP(f)+j} R(r, c) \quad (5-1)$$

In this expression, $R(r, c)$ is the radiance as a function of row r and color slice c , and $M_{f,c}$ corresponds to the number of rows contained in the footprint-color slice pair. The sum is over all rows r contained in footprint f (as defined in the last column) but shifted by j pixels. In other words, if in the last column, a footprint is comprised of pixels $\{p, \dots, p+19\}$, sums are performed over pixels $\{p+j, \dots, p+19+j\}$, for all j from 0 to the number of clocking jumps contained in the band. Thus, these mean radiances represent averages over all the possible sets of pixels that may comprise a footprint.

Anomalous (e.g., negative) values must be eliminated so that they do not skew the averages, since they would lead to potentially large errors in the corrections. Therefore, a heuristic filtering is performed in which “bad” pixels are replaced with an average of nearest neighbors (in row space). Pixels are marked “bad” if they satisfy at least one of three conditions: (1) they are marked as “bad” in the ARP, (2) they have a negative value, or (3) the relative spatial gradient exceeds some threshold and the absolute spatial variation is more than twice the noise level σ , defined by:

$$\sigma = \frac{R_{max}}{100} \sqrt{\frac{100R(\lambda)}{R_{max}} \eta_0^2 + \eta_1^2} \quad (5-2)$$

where $R(\lambda)$ is the radiance, R_{max} is an estimate of the maximum expected signal, and η_j are the signal-to-noise ratio (SNR) coefficients defined in the L1B file. The threshold is an adjustable parameter and taken to be a relative increase of 1.6 (or decrease of 1/1.6) from one row to the next. *The current scheme fails if there are two or more consecutive bad pixels.* This almost never happens for the O₂ A-band, but does occur occasionally in the SCO2 band.

Table 5-1 provides a glossary of symbols used in the clocking algorithm description.

Table 5-1. Glossary of symbols.

Symbol	Definition
$\bar{R}(f, c, j)$	Mean radiance, over rows $\{r + j \mid r \in f\}$ in footprint f and color slice c
$M_{f,c}$	Number of pixels in color slice c , footprint f
$R(f, c)$	Radiance at row r in color slice c
$r(f, G, j)$	Ratio needed to correct clocking at jump j for footprint f and group G
$N_{f,G}$	Normalization factor for footprint f and group G to preserve mean radiance over entire spectrum
$h_G(\lambda)$	Weighting function for group G at column λ
$C_G(f, \lambda)$	Correction function for group G , footprint f at column λ
$C(f, \lambda)$	Final correction function applied to L1B radiance for footprint f , column λ

Figure 5-10 shows a large but not extreme example of the spatial variation that might appear within a single frame. The grey, dotted vertical lines indicate footprint boundaries. Individual solid-colored curves correspond to the *corrected* radiance within particular color slices, while all uncorrected radiance measurements are denoted by grey circles. Footprints 3 and 4 (counting from the left and starting from zero) show significant differences in both radiance and slope between their left and right boundaries. Such differences will produce large clocking corrections. Note also that there are three uncorrected points whose nearly vanishing values are replaced by averages of their neighbors. Without this correction the mean radiances within a footprint would be underestimated by $\sim 5\%$. If the “bad” pixels occur near footprint boundaries and were to remain uncorrected, the clocking algorithm would produce erroneous corrections.




Figure 5-10. O₂ A band radiances measured for a few select color slices as a function of row index, which serves as a proxy for actual spatial variation within a scene. Solid curves are corrected values in which “bad” pixel data have been removed. Grey circles represent original measured data. For this particular case, three bad pixels have been replaced by nearest neighbor averages. Grey vertical lines represent footprint boundaries. The large differences between the values of the radiance on the boundaries of footprints 3 and 4 (starting from zero on the left) will produce large “clocking errors,” or spectral discontinuities in the measurement.

5.2.1.2 Second Stage: Computing Ratios and Final Correction Function

After computing and saving the radiance means, the algorithm proceeds to the second stage in which column radiances are modified. First, the set of color slices is partitioned into “groups”, which consist of clusters of spectrally close columns. The motivation for this partition is statistical. Using color slices that are spectrally very close should produce similar discontinuities; averaging over a small ensemble of strongly correlated slices reduces statistical error. The aggregation of color slices into groups follows a simple rule: two slices are in the same group if they are separated by fewer than n pixels, where n is taken to be 6 by default but is adjustable. (If n were chosen to be 0, for example, the groups would all be singletons.) In the future this rule may be further refined, possibly by allowing for aggregations that are dependent on footprint. Illustrations of the current selection of color slices and their aggregation into groups for each band are shown in Figure 5-5, Figure 5-7, and Figure 5-9.

We now combine the mean radiances computed in the first stage into a group dependent quantity by summing the mean radiances over all color slices contained in a group. Then, for each group, the ratio of the mean radiance at each jump is computed:

$$r(f, G, j) = \frac{1}{N_{f,G}} \frac{\sum_{c \in G} \bar{R}(f, c, j)}{\sum_{c \in G} \bar{R}(f, c, j + 1)} \quad (5-3)$$

In this expression the summation is over all color slices, c , within a group, G , and $N_{f,g}$ is a normalization factor defined below. (All radiance variables depend also on frame index, which has been omitted for notational clarity.) Neglecting the normalization factor for the moment, these ratios define multiplicative factors that if applied to only one side of a jump, produce a continuous radiance across the jump.

We now need to interpolate (and also extrapolate) to all spectral points. The ratio described above specifies a multiplicative factor to be applied at a jump position, p under the constraint that the multiplicative factor to be applied at position $p+1$ is identically equal to unity. To extend the multiplicative factor to all columns, we linearly interpolate from a value of one from the last clocking jump to the value given by the ratio as defined above. The resulting correction function, $C_G(f, \lambda)$, is a saw tooth defined for each group G and at every spectral point λ . The left panel of Figure 5-11 shows example correction functions for a particular footprint for the O₂ A-band. Plotted are several curves of slightly different shade of blue, corresponding to correction functions using different groups. In this example, all ratios at all jumps were found to lie between 0.95 and 0.97, and for each group a correction function is defined by linearly interpolating between a jump location, where the function's value is the ratio and one pixel to the right of the previous jump location, where the function's value equals one. To maintain roughly the same average radiance across the entire band, this correction function is normalized by a factor $N_{f,G}$ such that its mean over the entire spectral domain is equal to unity. The right panel of Figure 5-11 features red curves corresponding to the blue ones in the left panel, but after renormalization.




Figure 5-11. Example of a clocking correction function for the O₂ A band. The blue curves on the left correspond to different pre-normalized correction functions determined by different groups of color slices for a single frame and footprint. To (roughly) preserve the value of the average radiance, the correction functions are normalized such that the mean over columns equals one, as shown in the red curves on the right. Multiplication by weighting functions for each group yields the final correction function, depicted as the black curve on the right.

At this point, one has multiplicative correction functions for each color slice grouping. These functions need to be combined to produce a single correction function, $C(f, \lambda)$, for each frame-footprint pair. Because the corrections are expected to be most accurate for columns that are spectrally close to a group, we employ simple hat functions to create a weighted average:

$$C(f, \lambda) = \sum_G h_G(\lambda) C_G(f, \lambda) \quad (5-4)$$

where $h_G(\lambda)$ are the triangular weighting functions, which fall linearly from a maximum value of one at the average column position of group G to a value of zero at the average positions of the neighbors of G . If G is the first or last group—that is, it does not have a neighbor to the left or to the right—then $h_G(\lambda)$ has a constant value of one from the average column position to the end of the detector. Note that the hat functions depend only on the choice of color slices and their grouping and have the property that:

$$\sum_G h_G(\lambda) = 1 \quad (5-5)$$

The weighting functions for the current selection of color slices and groupings for the O₂ A-band is shown in Figure 5-12, where each color corresponds to a different group, G .

The sum of the product of these weighting functions and the curves in red in Figure 5-13 yields the weighted average correction function, shown in black in Figure 5-13. This weighted correction is applied to the original L1B radiances to produce the corrected radiances:

$$R_{corrected}(f, \lambda) = R(f, \lambda)C(f, \lambda) \quad (5-6)$$




Figure 5-12. Weighting functions for each group of color slices for the OCO-2 O₂ A band.

In practice, most corrections for non-lunar scenes are fairly small; the difference between the original and corrected radiances is barely perceptible, although discontinuities may be more evident in the residuals of retrievals. A typical “large” correction is on the order of several percent. To facilitate analysis of the effect of the clocking correction a simple scalar metric, $\mu(f)$, depending only on sounding is also saved to the L1B output:

$$\mu(f) = \max_{\lambda} |C(f, \lambda) - 1| \quad (5-7)$$

Figure 5-13 shows the effect of applying the correction illustrated in the preceding figures to the measured radiance of the O₂ A-band. The black/red curves denote original/corrected radiances. Figure 5-13 shows that a clocking jump occurs at column 843. The inset illustrates the elimination of the discontinuity at that pixel to produce an even continuum region about the O₂ doublet.




Figure 5-13. An example OCO-2 O₂ A spectrum, both before and after clocking correction. The inset shows a close up in the neighborhood of a jump, where the correction is largest.

5.2.2 Validation of the Clocking Correction

The behavior of the clocking correction varies significantly among the three bands. In this section, we restrict discussion to a comparison of the effect of the correction on a representative L1B product (orbit 1600); its impact on actual XCO₂ retrievals will not be considered here. Orbit 1600, shown in Figure 5-14 contains both land (33%) and ocean frames (66%).

Figure 5-15 shows the cumulative distribution of the magnitude of the correction (as defined above) for each band and terrain type (i.e., land and water). For example, for the O₂ A-band roughly 90% of all frame-footprint pairs over land require a correction of less than 0.5%. While the statistics for the O₂ A and WCO₂ bands are similar over land, the SCO₂ bands requires larger corrections: the 90th percentile for the SCO₂ occurs at a correction magnitude of 1%. The reason for this difference is unclear but may be related to the much larger number of “bad” pixels in the SCO₂ band. We note that for each case shown in Figure 5-15 at least 40% of scenes require no correction at all. This result is an artifact of the noise cutoff threshold discussed in the previous section; corrections that are of the same order of magnitude as the noise are artificially set to zero.

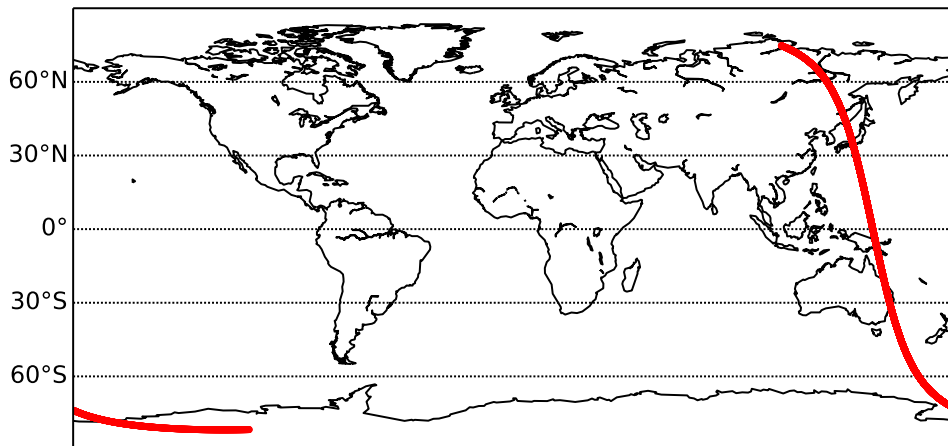


Figure 5-14. Orbit used in discussion of the clocking correction

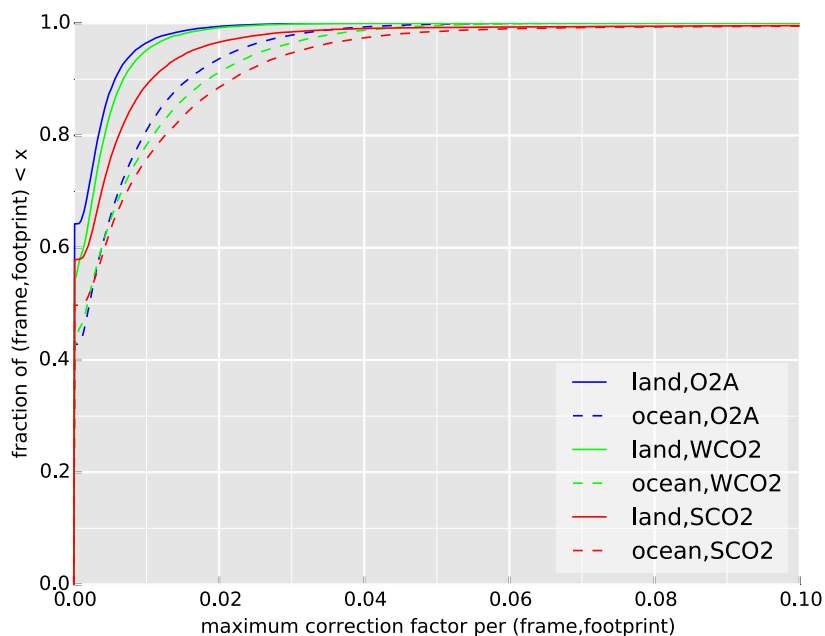


Figure 5-15. Cumulative distribution function of magnitude of the correction for orbit 1600. Results are separated by band and terrain type (i.e., land and water).

One useful metric to assess the soundness of the correction of the algorithm is to examine the consistency of the predicted corrections among the color slices. Figure 5-13, Figure 5-17, and Figure 5-18 show for the O₂ A, WCO₂, and SCO₂ bands, respectively, the correlation coefficient, ρ , of the ratios r , defined in the previous section:

$$\rho(f, G, G') = \frac{\sum_{s,j} r(s, f, G, j)r(s, f, G', j)}{(\sum_{s,j} r^2(s, f, G, j))^{1/2}(\sum_{s,j} r^2(s, f, G', j))^{1/2}} \quad (5-8)$$

where the summation over a frame index variable s has been made explicit and instead of the employing the grouping described previously, we have taken each group G to contain a single

color slice. It was shown previously that a significant fraction of corrections is either small or identically zero. Since large relative differences of very small quantities may not be meaningful, we impose a restriction to the summation in the equation above: for a given group (i.e., color slice) G we sum over only those corrections that are larger than a cutoff of 0.5%. Note that because the set of summands for a group G will not be identical to the set of summands for a group G' , the correlation coefficient defined above will not generally be a symmetric function of G and G' .




Figure 5-16. Correction correlation matrix for the O₂ A band.

Figure 5-16 shows the correlation matrix by footprint for the O₂ A-band. We first note that the correlations are rather strong in all cases, everywhere exceeding 0.87. Also, as expected, color slices that are spectrally close have stronger correction correlations. This finding justifies our aggregation of color slices into groups. However, note that there is a dependence on footprint: spectrally distant color slices show greater consistency for lower footprint indices than for higher indices. That is, for higher footprint indices a more diverse sampling of color slices is required for the correction algorithm. This result is most apparent for the O₂ A-band but is evident for the other bands as well. Finally, it was noted earlier that slices 16, 17, and 18 (corresponding to columns 68, 67, and 66) actually lie near solar lines and therefore ought not to have been included in the correction algorithm. It was asserted that their inclusion, nonetheless, would not significantly modify the final outcome. The strong correlation illustrated in Figure 5-16 between those color slices and all the other slices on the shortwave end justifies that assertion.

Figure 5-17 shows correlations between color slices for the WCO2 band. Generally, the correlation matrix elements are even stronger than for the O₂ A-band even for spectrally distant slices. This is likely due to the low degree of clocking in the WCO2 band, which suffers only three jumps over the whole band as compared to four in the O₂ A-band and eight for the SCO2 band. However, there are some oddly weaker correlations involving color 211 in footprints 5 and 6 and color 922 in footprints 0 and 1. We suspect these arise from a greater number of instances where consecutive “bad” pixels foil the correction algorithm.




Figure 5-17. Correction correlation matrix for the WCO2 band.

Unlike the other two bands, the corrections prescribed by color slices in the SCO2 band are not generally very consistent, even for spectrally nearby slices as Figure 5-18 shows. We strongly suspect that this result stems from the much larger number consecutive “bad” pixels, which the algorithm fails to properly handle. Note that the group of color slices 1–5 (colors 89 through 85) that do generally produce consistent corrections lie on the shortwave end of the detector, where there are far fewer pixels that are marked as “bad” in the ARP. Furthermore, the large clocking effect in the SCO2 should weaken the correlation between the two groups on the long and short wavelength ends of the detector.




Figure 5-18. Correction correlation matrix for the SCO2 band.

Finally, we consider the effect of the algorithm on the spectra themselves. One method to visualize the reduction of clocking-induced discontinuities is to perform a singular value decomposition on a set of spectra before and after correction and examine the resulting change in the right singular vectors:

$$S = U\Sigma V^* \quad (5-9)$$

In this expression the columns of S are comprised of the uncorrected spectra, Σ is the diagonal matrix of singular values, and the columns of U and V are the left and right singular vectors. Figure 5-19 demonstrates the effect of the clocking algorithm on the set of spectra in orbit 1600 for the O₂ A-band. The first eight singular vectors are plotted both before (top row) and after (bottom row) correction. Discontinuities that are clearly visible in the sixth (yellow-green), seventh (orange), and eighth (red) singular vectors prior to correction are completely eliminated by the clocking algorithm. Because only a minority of spectra require significant correction, the size of the effect as measured by the ratio of singular values is fairly small; the ratio of sixth and first singular value is 0.0045 in this data set. The impact of the correction on the other bands is similar and shown in Figure 5-20 and Figure 5-21.

These results indicate that this correction approach will usually be adequate. In cases where it is not, it will produce a poor fit in the Level 2 retrieval, and the Level 2 post processing screening algorithm will usually reject the sounding.




Figure 5-19. Right singular vectors corresponding to eight largest singular values from violet (largest) to red (smallest) for the O₂ A-band. The top/bottom rows are pre/post corrected vectors. The panels on the right are close up of clocking jumps at 843. The right singular vector corresponding to the sixth (yellow-green), seventh (orange), and eighth (red) largest singular values exhibit discontinuities at pixel 843 prior to correction.




Figure 5-20. First eight right singular vectors for the WCO₂ band. The clocking jump evident in the sixth one (in yellow-green) with singular value 0.28% of the largest value vanishes after the clocking correction.




Figure 5-21. First twelve right singular vectors for the WCO2 band. The clocking jump evident in the eleventh one (in orange) with singular value 0.34% of the largest value vanishes after the clocking correction.

6 Polarimetric Calibration

6.1 Treatment of Polarization

As noted above, the OCO-2 and OCO-3 instruments measure a single linear component of the polarization of the incoming light, and the polarization acceptance axis is perpendicular to the long axis of the spectral slit. The axis of polarization accepted by the instrument depends on various instrument and spacecraft geometry parameters. For convenience, this is given in the L1B file via two sets of parameters. The first is the polarization angle of the axis of the accepted polarization, and the second is a set of “Stokes Coefficients” $\{m_1, m_2, m_3, m_4\}$ which define the response of the instrument to Stokes vector of the incoming light:

$$I_{\text{meas}} = m_1 I + m_2 Q + m_3 U + m_4 V. \quad (6-1)$$

The polarization angle ϕ_{pol} is defined such that $m_2 = \frac{1}{2} \cos(2 \phi_{\text{pol}})$ and $m_3 = \frac{1}{2} \sin(2 \phi_{\text{pol}})$, and physically represents the angle between the axis of accepted polarization of the instrument and the reference plane for polarization. Stokes parameters Q and U also depend upon the chosen reference plane for polarization. We choose the local meridian plane to be this reference plane for polarization, which is the plane containing the local normal unit vector and the vector pointing from the target FOV to the satellite (Figure 6-1). Stokes parameter Q is then defined as:

$$Q = I_{\parallel} - I_{\perp}; \quad (6-2)$$

i.e., Q is the difference in intensities of the light polarized parallel and perpendicular to this plane. U is similar but requires a rotation by 45° :

$$U = I_{45} - I_{135}. \quad (6-3)$$

The definition of Stokes parameter U requires a definition of the “sense of rotation.” Excerpting from Schutgens et al. (2004):

“The sense of rotation is usually defined by demanding that the three unit vectors \perp , \parallel , and prop form a right-hand coordinate frame, see Figure 2. Chandrasekhar [1960], van de Hulst [1981], and Hovenier and de Haan [1985] assumed $\perp \times \parallel = \text{prop}$, while Azzam and Bashara [1987] and Slijkhuis [2000] assumed $\parallel \times \perp = \text{prop}$. The difference in Stokes vector is the sign of U (and hence χ $180^\circ - \chi$). Note that POLDER data [Bréon and CNES Project Team, 1997] uses the latter convention, while SCIAMACHY data [Slijkhuis, 1998] currently uses the first.”

Here, “prop” denotes a unit vector pointing in the direction of propagation. We use the former convention of Chandrasekhar (1960), van de Hulst (1981), and Hovenier and de Haan (1985), also adopted by the SCIAMACHY team (Slijkhuis 1998), i.e., we assume that $\perp \times \parallel = \text{prop}$. This convention is used both in the definition of the polarization angle and Stokes coefficients provided in the L1B files, as well as the L2 radiative transfer solvers.




Figure 6-1. Geometry of the reference plane for polarization, i.e., the local meridian plane. This plane is formed by the local normal unit vector and the ray from the target FOV to the satellite. The direction of propagation of measured light is the latter. Two example polarizations are shown depicting light polarized both parallel (\parallel) and perpendicular (\perp) to this plane. The polarization angle of the instrument is defined as the angle between the axis of accepted polarization and the reference plane for polarization.

Even in nadir mode, the target field of view never sits directly at the subsatellite point. It is always offset slightly, and there is a different geometry for each of the eight footprints in a given frame. Therefore, the polarization reference plane defined above can be significantly different for each of the eight footprints in a given frame, and therefore m_2 and m_3 may vary significantly from one footprint to the next in a given frame.

In previous documents, we indicated that the polarization reference plane would be defined as the principal plane, which is formed by the ray from the target to the Sun and the target to the satellite. This frame, while simpler in the sense that the definitions of Q and U would not vary much in nadir mode from one footprint to the next, is typically not used in the RT community. Therefore, we have avoided it in favor of the local meridian plane. In glint mode, the principal and local meridian planes are nearly identical, but in other observation modes they can differ significantly.

6.2 OCO-2 Polarization

For OCO-2, the orientation of the spectral slit, and thus the polarization angle, is controlled by the attitude of the spacecraft. The intent for OCO-2 was to have the orientation of the axis of accepted polarization horizontal to the ground for glint mode observation (i.e., perpendicular to the local meridian plane). A problem with the actual polarization plane of the instrument was discovered on orbit, and also found in OCO-3 during preflight testing. A mitigation strategy was employed to yaw the OCO-2 spacecraft by 30° in glint and nadir modes. In nadir mode, the polarization angle can vary widely because the meridian planes for each footprint can be significantly different as discussed above. In target mode, because of the complicated slewing of the spacecraft, the polarization angle can vary widely.

In the special case of glint measurements in the principal plane, when $\phi_{\text{pol}} = -30^\circ$, the measured intensity reduces to:

$$I_{\parallel} = 0.5 I + 0.25 Q - 0.433 U . \quad (6-4)$$

However, in nadir and target modes, the polarization angle may vary with the footprint index, for any value between 0° and 360° .

6.3 OCO-3 Pointing Mirror and Polarization

A design change that impacts the radiometric characteristics of the data is the addition of a pointing mirror assembly (PMA). The PMA is required to allow non-nadir observations from the fixed position on the ISS, unlike OCO-2, which maneuvers the entire spacecraft to point. Two important design requirements of the PMA were to allow quick movement through a large range of angles, and that the movement not impart any angular dependent polarization or radiance changes in the measurements. To meet these objectives a variation of the pointing system designed for the Glory Aerosol Polarimetry Sensor (APS) (Persh et al. 2010) was selected. This system relies on a single pair of matched mirrors in an orthogonal configuration that impart less than 0.05% change to the polarization (Mishchenko et al. 2007). For the OCO-3 PMA the concept was extended to a 2-axis pointing system. Although the PMA itself does not change the polarization of the light by more than 0.1%, there are still polarization implications because the image of the slit is rotated as a function of the change in the PMA position. It is worth noting that reflected sunlight is naturally polarized by its interaction with the Earth's surface and atmosphere, especially over water.

Early in the mission design, a trade study was performed evaluating the expected signal with and without the installation of an additional polarization scrambler, i.e., a polarization nulling optical component. Without a scrambler, the signal measured by the instrument is large when the polarization of the incident light is aligned with the axis of polarization sensitivity of the instrument. However, the measured signal decreases towards zero as the polarization of the incoming light rotates orthogonal to the axis of instrument polarization. Inserting a scrambler would make the polarization orientation of the incoming light random, regardless of the position of the PMA, but would reduce the signal by nearly 50% compared to having no scrambler. In addition, the analysis showed that a single optical element that could scramble light at all of the OCO-3 wavelengths could not be manufactured and characterized to the required precision. Lastly, the volume and coverage of data with sufficient signal in the no-scrambler case was predicted to be more than sufficient to meet the science objectives. Therefore, OCO-3 operates without a polarization scrambler. The PMA is one element that contributes to a change in the overall light throughput of OCO-3 compared to OCO-2. In the O_2 A-band, each mirror has a reflectivity of 95.4%, so the 4-mirror PMA system has an effective transmission of 83%. The weak and strong CO_2 band overall transmissions are higher, at 93% and 95%, respectively.

6.3.1 PMA Movement in Non-nadir Modes

During non-nadir observations with OCO-3 (glint, target, or snapshot area maps), the PMA will be regularly repositioned in order to follow the desired target location as the ISS moves through its orbit. The two-axis PMA angular position is controlled by two actuators, and their positions are reported within the SoundingGeometry data in the L1B files (pma_azimuth and pma_elevation). Note the elevation actuator angle is defined such that 0 degrees places the instrument boresight at the nadir direction from the ISS. The elevation actuator can move to positive and negative angles, and the azimuth actuator rotates the PMA around the nadir vector.

In practice, because of small variations in the precise ISS attitude and the orientation of the instrument platform, zero PMA elevation will be a few degrees away from nadir.

The full range of azimuth angles is limited due to space and hardware constraints in the instrument's installed location on the ISS. About half of the full azimuth travel is inaccessible. Because the elevation actuator has travel range to both positive and negative elevations, the entire viewing hemisphere can be reached. A view direction at a particular azimuth and elevation angle can also be reached by shifting the sign of the elevation angle, and shifting the azimuth by 180 degrees. Mathematically, for azimuth ϕ and elevation θ , (ϕ, θ) is equivalent to $(\phi + 180^\circ, -\theta)$. However, since the polarization angle depends on the actual angle reported by the PMA elevation actuator (including its sign), we do not convert the angles in this way, but instead use the actuator positions as reported.

An example path in the PMA actuator angle space is shown in Figure 6-2 using data from a real early mission orbit. To correctly display the positive and negative elevations, we split the orbit into sections for the two elevation signs. For positive PMA elevation actuator angles (Figure 6-2, left), the boresight will be aimed forward of the ISS (in other words, ahead of the ISS, along the satellite track) when the PMA azimuth actuator is at 270° , and aimed to starboard at 0° azimuth. Negative PMA elevation (Figure 6-2, right) reverses these directions, such that the boresight is aimed backward (behind the ISS, along the satellite track) at 270° azimuth and aimed to port at 0° azimuth. The figure illustrates a typical path that the PMA takes during ocean glint observations, and in this case, the majority of the orbit consisted of glint observations. The first portion of the orbit starts with the PMA at the maximum elevation (as the glint spot will appear closest to the horizon at the start of the collection), and is nearer toward the 0° azimuth, indicating the glint spot is to the starboard side of the ISS. The PMA moves towards smaller elevation while also increasing azimuth and moving toward the inaccessible region (gray region on right). When the PMA nears this region, it will flip by nearly 180° in azimuth, and switch to negative elevation, to continue following the path of the glint spot. In negative elevation, the PMA azimuth will be nearer to the 180° azimuth position to follow the glint spot on the starboard side.

Portions of the hemisphere are defined as “keep out zones”, where parts of the ISS (the solar panels in particular) could be in the instrument field of view. The inflight control software limits the PMA such that it cannot be commanded to point the instrument in these directions. The keep out zones are not labeled in Figure 6-2 and are not explicitly considered in the description of the PMA in this document.




Figure 6-2. Polar coordinate view of the PMA actuator angles, with the PMA elevation plotted as the radial coordinate, for positive elevation (left) and negative elevation (right). The gray region on the left side of each plot shows the range of azimuth angles that are inaccessible due to physical limitations. The plot points show the PMA angular locations for part of a real OCO-3 data collection where the instrument was collecting glint mode data over ocean.

6.3.2 PMA Polarization Angle

As the PMA actuators move, the orientation of the entrance aperture will rotate around the boresight. Since the polarization acceptance angle is directly related to the aperture orientation, this implies the polarization angle will also change as the PMA actuators move. Physically, movement in the azimuth actuator rotates the aperture an equal amount counterclockwise (so that a movement of x degrees in PMA azimuth will also change the polarization angle by x degrees), and movement in the elevation actuator moves the aperture by an equal amount clockwise (a movement of x degrees in PMA elevation will change the polarization angle by $-x$ degrees). By assuming an idealized model for the sensor to earth geometry where the sensor is aligned so that PMA elevation = 0 is aligned to the subsatellite nadir vector, and assuming a spherical Earth, then the polarization angle can be expressed as the following simple equation, with the polarization angle (χ) related to the PMA azimuth (φ_{PMA}) and PMA elevation (θ_{PMA})

$$\chi = \begin{cases} 90.0 - \theta_{\text{PMA}} + \varphi_{\text{PMA}} & \text{if } \theta_{\text{PMA}} > 0 \\ 270.0 - \theta_{\text{PMA}} + \varphi_{\text{PMA}} & \text{if } \theta_{\text{PMA}} < 0 \end{cases} \quad (6-5)$$

In practice, this simple equation will not account for deviations from the idealized model, nor with the adjustments needed across the slit position for the eight spatial footprints. Figure 6-3 shows the difference between the simple model and polarization angles computed from the full operational calculations that use the Earth geoid, the instrument geometric model and the real ISS attitude. Throughout the following sections, we use the simplified expression in the equation when computing the expected polarization response for the simulated data.

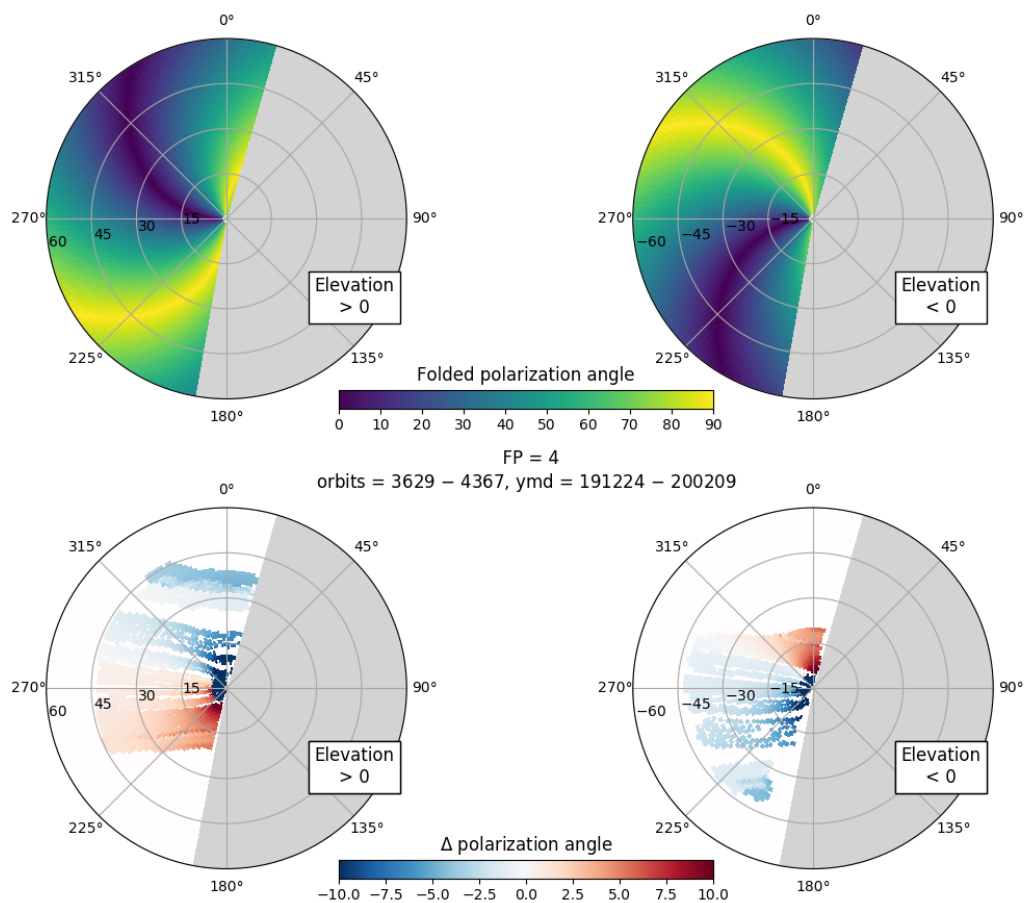


Figure 6-3. Top: View of the polarization angle (using the approximation in the equation above), folded to the range 0–90 degrees, using the dual polar coordinate axis (same as Figure 6-2). For ocean glint observations, the view geometry is in the principal scattering plane, and the incident radiance will tend to be polarized at an angle of 90 degrees. Therefore, regions with small polarization angles (dark blue colors) will have the lowest SNR. Bottom: The difference between the approximate equation and the exact calculation using the full geolocation algorithm, using several weeks of early mission ocean glint data from OCO-3.

6.3.3 PMA Polarization Angle and ISS β Angle

For glint ocean observations, the PMA must move through particular curved paths in azimuth-elevation space in order to track the relative movement of the glint spot throughout the ISS orbit. Figure 6-2 showed a path for a particular orbit. The PMA path through the angle space is primarily dependent on the ISS orbital β angle, which is the angle between the ISS orbital plane and the vector from the Earth to the Sun. Figure 6-4 shows a set of paths at a particular range of ISS β angle from the real ISS ephemeris from 2015 (Eldering et al. 2019). For each data subset, the plotted data consist of all points in the simulation within a range of 1° centered at the listed β angle (for example, the data displayed at $+60^\circ$ is all points in the range 59.5° – 60.5°). Thus, each grouping of data at a particular β value will sample multiple different times of year. The width of the data distribution within each β range is driven by smaller scale variations to the ISS altitude, orbital inclination, and the variation of the Earth geoid due to different latitudinal sampling.

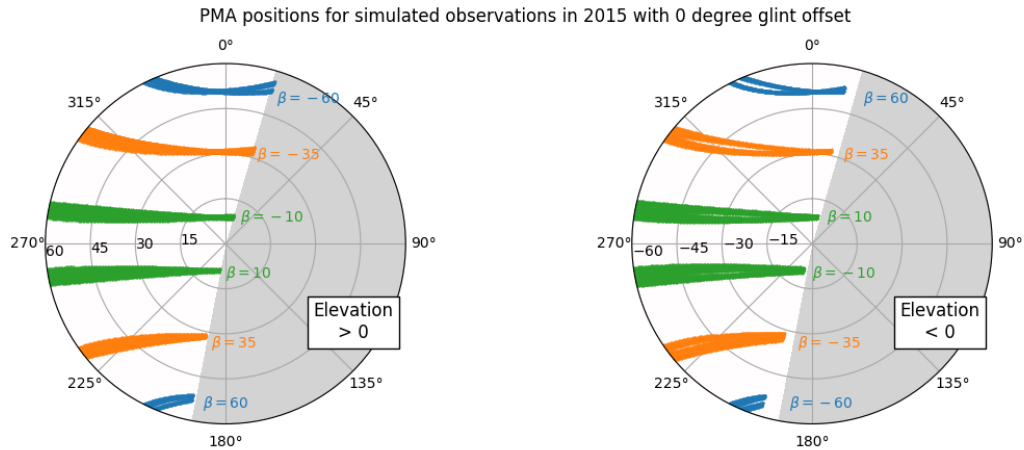


Figure 6-4. Example paths in the PMA actuator angle space for glint observations, at a range of ISS β angles.

The range of paths shown in Figure 6-4 must then be compared to the polarization angles shown in in Figure 6-3. Certain β angles require paths in PMA angle space that will rotate the boresight into unfavorable polarization angles (primarily negative values of β). The impact of the polarization angle on the expected SNR depends on the relative amount of polarization in the observed radiance, which is related to the incident angle of the radiation (the solar zenith angle). The reflected radiance has very little polarization at small solar zenith angles (less than about 20°), maximum polarization at the Brewster angle (about 53° for water), and decreasing again toward larger solar zenith angles. Similar to Figure 6 in Elderling et al. (2019) we can summarize the SNR relationships by plotting the paths through the polarization angle and solar zenith angle parameter space, as displayed in Figure 6-5. A theoretical SNR prediction is displayed with the background contours. This prediction uses a Cox-Munk polarized reflection model with a fixed wind speed and zero offset angle between the boresight and the glint spot. Due to the symmetries of the polarization angle in this context, only the angles 0° – 90° need to be considered (the remainder of the domain will contain copies of this pattern). The paths taken through this parameter space are directly related to the ISS β angle as before, and the paths near -40° stand out as the most impacted by unfavorable polarization angle. For those values of β angle, the PMA will move through the 0° and 180° polarization angles near where the viewing geometry has the incident light close to the Brewster angle. The up and down triangle symbols denote the positive and negative PMA elevation segments of the total observation path. The two paths are nearly symmetric, but not entirely, due to slight geometry differences between the PMA elevation signs.

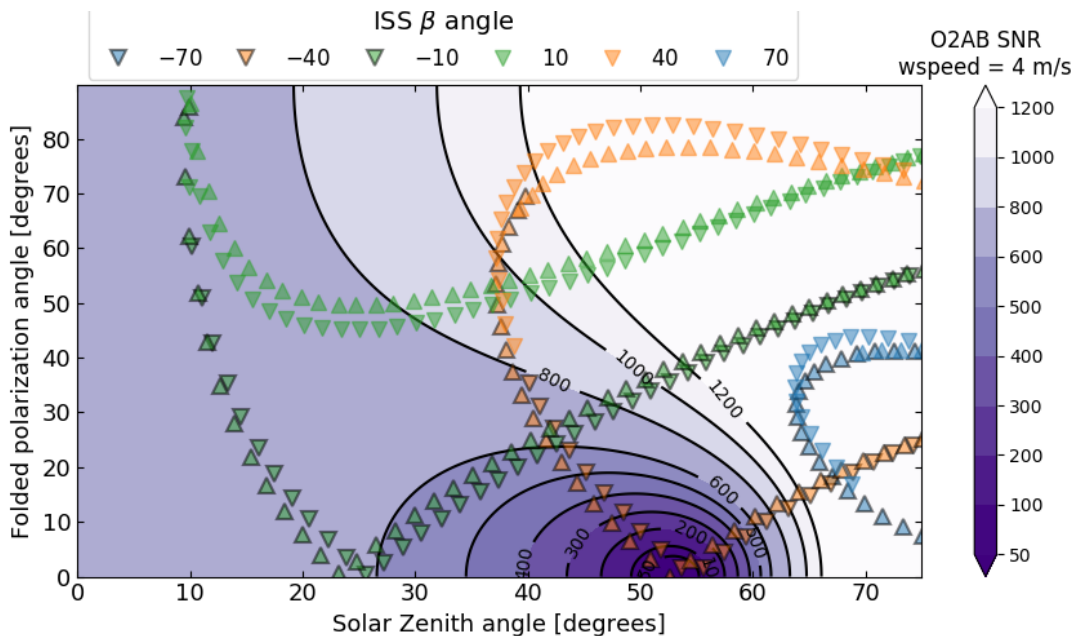


Figure 6-5. SNR calculation for Cox-Munk surface model as a function of Solar Zenith angle and polarization angle (background filled contours), with over plotted points for various paths followed by the PMA for glint mode observations. The over plotted symbols are colored by the absolute value of ISS β (70 = blue; 40 = green; 10 = orange), with the negative ISS β symbols including a dark gray border. The triangles with a up or down pointing refer to positive or negative PMA elevation angles.

7 Cosmic Ray Artifacts

Cosmic rays rarely produce permanent damage to the OCO-2 FPAs, but energetic particles produce ion trails as they traverse the FPAs, producing spurious intensity spikes. The largest effects are seen in the O₂ A-band, because the FPA in this channel is more sensitive to cosmic ray artifacts. Primary and secondary cosmic rays with a broad range of energies are occasionally seen everywhere along the orbit path but they are most common in the vicinity of the South Atlantic Anomaly (SAA), where up to 2% of an O₂ A-band spectral range can be contaminated in individual frame. The approximate geographical extent of the SAA as it impacts the OCO-2 mission can be seen in Figure 7-1. The characteristic spectral signature of cosmic rays is shown in Figure 7-2.

An algorithm has been developed to identify and screen these events in this L1B data delivered to the science community. This algorithm includes the following steps:

1. For each L1 granule, a singular value decomposition is performed independently for each footprint and spectral band.

L2/L1bScSpectralParameters/spike_eof_bad_colors_o2

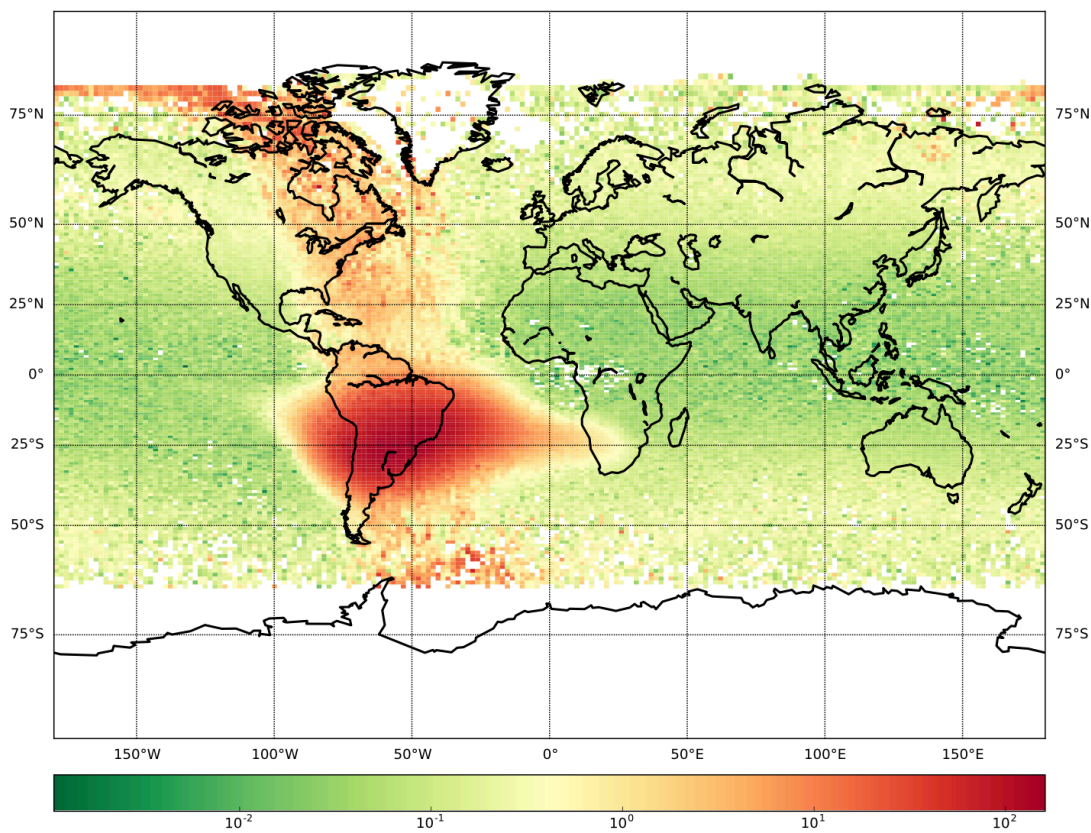


Figure 7-1. The number of spectral samples in the OCO-2 O₂ A band that are not contaminated by 10-sigma cosmic ray events clearly shows the geographical extent of the South Atlantic Anomaly (SAA).




Figure 7-2. An example showing the impact of a radiation event on spectra. Note the three spikes with unusually high intensities.

2. Using the 40 leading eigenvectors, a least squares approach fits the optimal linear combination of all 40 eigenvectors to match individual spectra within the granule. This represents an empirical fit to all spectra at a computational speed that is not prohibitive for L1B processing. For strong outliers (>10 sigma), spectral samples are flagged and refitted after excluding outliers that strongly perturb the linear fit.
3. The fit mismatch (residual), weighted by the expected detector noise is calculated “(measurement-model)/noise” and saved for each spectral sample.

In the L1B files these weighted residuals are saved in the group SpikeEOF as rounded integer values with all residuals less than 3 three times the noise level (3 sigma) set to 0 to minimize file size. The variables used to record the weighted residuals in the L1B files are:

SpikeEOF/spike_eof_weighted_residual_band (with *band* = o2, strong_co2, weak_co2).

The total number of spectral samples contaminated in by cosmic rays in any one spectrum are recorded in the variables:

SpikeEOF/spike_eof_bad_colors_band (with *band* = o2, strong_co2, weak_co2).

The OCO-2 L2 FP algorithm uses the weighted residuals in the L1B file to flag individual spectral samples as bad if their weighted residuals exceeds +6 sigma, and if the sounding falls within a geographic boundary box around the SAA (latitude -50 to 0 , longitude -90 to 10 degrees). We use a one-sided threshold value as cosmic ray hits can only cause a positive anomaly in measured spectra. L1B users are advised adopt a similar approach for flagging spectral samples contaminated by cosmic rays before attempting any fit. We also advise users to apply this threshold only in the vicinity of the SAA. The same approach is being applied to OCO-3.

8 Ancillary Radiometric Product (ARP) Files

The key calibration parameters are stored in the ARP, which is not distributed publicly. These files are changed as needed and deemed necessary, roughly weekly. There will be an ARP file for forward processing and an updated ARP for reprocessing to reflect any additional knowledge gained from the measured data set about dark subtraction or gain degradation. Each data product identifies the ARP used in Metadata/ARPAncillaryDatasetDescriptor.

9 Other Resources

There are a number of other project documents that the user should refer to as they work with the data.

- The OCO-2/OCO-3 Level 2 Data User’s Guide—this document discussed the publicly distributed data products, L1bSc, L2IDP, L2Std, and L2Dia. The key data fields are discussed, and tables are included that specify all of the fields in each data product.
- The OCO-2/OCO-3 L2 ATBD—this ATBD steps through the physics and implementation of the Level 2 algorithm.
- The ATBDs for the IMAP-DOAS and ABO2 algorithms—these ATBDs describe the two methods of identifying potentially cloudy footprints, in what we refer to as the prescreening step. These data are then used for setting data quality and data selection levels.
- Published papers—there are a number of published papers describing the algorithm, application to GOSAT, prescreening steps, etc. Please see the references below or the most up-to-date list of publications at

<https://ocov2.jpl.nasa.gov/science/publications/>

<https://ocov3.jpl.nasa.gov/science/publications/>.

10 References

- Bruegge, C. J., Crisp, D., Helmlinger, M. C., Kataoka, F., Kuze, A., Lee, R. A., McDuffie, J. L., Rosenberg, R. A., Schwandner, F. M., Shiomi, K., and Yu, S. (2019). Vicarious calibration of orbiting carbon observatory-2. *IEEE Transactions on Geoscience and Remote Sensing*, 57, 5135–5145. doi:10.1109/TGRS.2019.2897068.
- Chandrasekhar, S. (1960). Radiative Transfer. Dover Publications, Dover, N.Y.
- Crisp, D., Miller, C. E., DeCola, P. L. (2008). NASA Orbiting Carbon Observatory: measuring the column averaged carbon dioxide mole fraction from space. *J. Appl. Remote Sens.*, 2, 023508. doi:10.1117/1.2898457.
- Crisp, D. (2008). The Orbiting Carbon Observatory: NASA’s first dedicated carbon dioxide mission. In *Sensors, Systems, and Next-Generation Satellites XII*, 7106, 33–39. *SPIE*. doi: 10.1117/12.802194.
- Crisp, D., et al. (2017). The on-orbit performance of the Orbiting Carbon Observatory-2 (OCO-2) instrument and its radiometrically calibrated products, *in preparation*.
- Day, J. O., O’Dell, C.W., Pollock, R., Bruegge, C. J., Rider, D., Crisp, D., and Miller, C. E. (2011). Preflight spectral calibration of the Orbiting Carbon Observatory. *IEEE Trans. Geosci. Remote Sens.*, 49, 2793–2801.
- Eldering, A. (2019). The OCO-3 mission: measurement objectives and expected performance based on 1 year of simulated data. *Atmos. Meas. Tech.*, 12, 11400–11406.
- Frankenberg, C., Pollock, R., Lee, R. A. M., Rosenberg, R., Blavier, J.-F., Crisp, D., O’Dell, C. W., Osterman, G. B., Roehl, C., Wennberg, P. O., and Wunch, D. (2014). The Orbiting Carbon Observatory (OCO-2): spectrometer performance evaluation using pre-launch direct sun measurements. *Atmos. Meas. Tech.*, 7, 1–10. www.atmos-meas-tech.net/7/1/2014/doi:10.5194/amt-7-1-2014.
- Hovenier, J. W., and de Haan, J. F. (1985). Polarized light in planetary atmospheres for perpendicular directions. *Astron. Astrophys.*, 146, 185–191.
- Kuze, A., et al. (2013). Long-term vicarious calibration of GOSAT short-wave sensors: techniques for error reduction and new estimates of radiometric degradation factors. *IEEE Transactions on Geoscience and Remote Sensing*, 52.7: 1–14.
- Kuze, A., et al. (2011). Vicarious calibration of the GOSAT sensors using the Railroad Valley desert playa. *IEEE Transactions on Geoscience and Remote Sensing*, 49.5: 1781–1795.
- Lee, R. A. M., et al. (2017). Preflight spectral calibration of the Orbiting Carbon Observatory-2. *IEEE Transactions on Geoscience and Remote Sensing*, 55, No. 4, 2499–2508.
- O’Dell, C. W., Day, J. O., Pollock, R., Bruegge, C. J., O’Brien, D. M., Castano, R., Tkatcheva, I., Miller, C. E., and Crisp, D. (2011). Preflight radiometric calibration of the Orbiting Carbon Observatory, *IEEE Transactions on Geoscience and Remote Sensing*, 49.6: 2438–2447. doi: http://dx.doi.org/10.1109/TGRS.2010.2090887.

- Rosenberg, R., et al. (2017). Preflight radiometric calibration of Orbiting Carbon Observatory 2, *IEEE Transactions on Geoscience and Remote Sensing*, 55, No. 4, 1994–2006.
- Rosenberg, R., et al. (2020). OCO-2 calibration refinement across versions and plans for OCO-3, *IEEE International Geoscience and Remote Sensing Symposium*, 6381–6384
- Schutgens, A. J., Tilstra, L. G., Stammes, P., and Bréon, F.-M. (2004), On the relationship between Stokes parameters Q and U of atmospheric ultraviolet/visible/near-infrared radiation, *J. Geophys. Res.*, 109 (D09205). doi:10.1029/2003JD004081.
- Slijkhuis, S. (1998). SCIAMACHY level 0 to 1c processing algorithm theoretical basis document, Tech. Note ENV-ATB-DLR-SCIA-0041, Dtsch. Zent. für Luft- und Raumfahrt, Oberpfaffenhofen, Germany.
- van de Hulst, H. C. (1981), *Light Scattering by Small Particles*, Dover, Mineola, N. Y.
- Yoshida, Y., Kikuchi, N., and Yokota, T. (2012). On-orbit radiometric calibration of SWIR bands of TANSO-FTS onboard GOSAT. *Atmospheric Measurement Techniques Discussions*, 5.4: 4711–4734.