



## The on-orbit performance of the Orbiting Carbon Observatory-2 (OCO-2) instrument and its radiometrically calibrated products

David Crisp<sup>1</sup>, Harold R. Pollock<sup>1</sup>, Robert Rosenberg<sup>1</sup>, Lars Chapsky<sup>1</sup>, Richard A. M. Lee<sup>1</sup>, Fabiano A. Oyafuso<sup>1</sup>, Christian Frankenberg<sup>1,2</sup>, Christopher W. O'Dell<sup>3</sup>, Carol J. Bruegge<sup>1</sup>, Gary B. Doran<sup>1</sup>, Annmarie Eldering<sup>1</sup>, Brendan M. Fisher<sup>1</sup>, Dejian Fu<sup>1</sup>, Michael R. Gunson<sup>1</sup>, Lukas Mandrake<sup>1</sup>, Gregory B. Osterman<sup>1</sup>, Florian M. Schwandner<sup>1,4</sup>, Kang Sun<sup>5</sup>, Tommy E. Taylor<sup>2</sup>, Paul O. Wennberg<sup>2</sup>, and Debra Wunch<sup>2,6</sup>

<sup>1</sup>Jet Propulsion Laboratory/California Institute of Technology, Pasadena, CA, USA

<sup>2</sup>California Institute of Technology, Pasadena, CA, USA

<sup>3</sup>Colorado State University, Fort Collins, CO, USA

<sup>4</sup>Joint Institute for Regional Earth System Science and Engineering, University of California Los Angeles, Los Angeles, CA, USA

<sup>5</sup>Harvard-Smithsonian Center for Astrophysics, Cambridge, MA, USA

<sup>6</sup>Department of Physics, University of Toronto, Toronto, Canada

Correspondence to: David Crisp (david.crisp@jpl.nasa.gov)

Received: 28 August 2016 – Published in Atmos. Meas. Tech. Discuss.: 4 October 2016

Revised: 1 December 2016 – Accepted: 1 December 2016 – Published: 5 January 2017

**Abstract.** The Orbiting Carbon Observatory-2 (OCO-2) carries and points a three-channel imaging grating spectrometer designed to collect high-resolution, co-boresighted spectra of reflected sunlight within the molecular oxygen (O<sub>2</sub>) A-band at 0.765 microns and the carbon dioxide (CO<sub>2</sub>) bands at 1.61 and 2.06 microns. These measurements are calibrated and then combined into soundings that are analyzed to retrieve spatially resolved estimates of the column-averaged CO<sub>2</sub> dry-air mole fraction, XCO<sub>2</sub>. Variations of XCO<sub>2</sub> in space and time are then analyzed in the context of the atmospheric transport to quantify surface sources and sinks of CO<sub>2</sub>. This is a particularly challenging remote-sensing observation because all but the largest emission sources and natural absorbers produce only small (<0.25 %) changes in the background XCO<sub>2</sub> field. High measurement precision is therefore essential to resolve these small variations, and high accuracy is needed because small biases in the retrieved XCO<sub>2</sub> distribution could be misinterpreted as evidence for CO<sub>2</sub> fluxes.

To meet its demanding measurement requirements, each OCO-2 spectrometer channel collects 24 spectra s<sup>-1</sup> across a narrow (<10 km) swath as the observatory flies over the sunlit hemisphere, yielding almost 1 million soundings each day. On monthly timescales, between 7 and 12 % of these

soundings pass the cloud screens and other data quality filters to yield full-column estimates of XCO<sub>2</sub>. Each of these soundings has an unprecedented combination of spatial resolution (<3 km<sup>2</sup>/sounding), spectral resolving power ( $\lambda/\Delta\lambda > 17\,000$ ), dynamic range ( $\sim 10^4$ ), and sensitivity (continuum signal-to-noise ratio >400).

The OCO-2 instrument performance was extensively characterized and calibrated prior to launch. In general, the instrument has performed as expected during its first 18 months in orbit. However, ongoing calibration and science analysis activities have revealed a number of subtle radiometric and spectroscopic challenges that affect the yield and quality of the OCO-2 data products. These issues include increased numbers of bad pixels, transient artifacts introduced by cosmic rays, radiance discontinuities for spatially non-uniform scenes, a misunderstanding of the instrument polarization orientation, and time-dependent changes in the throughput of the oxygen A-band channel. Here, we describe the OCO-2 instrument, its data products, and its on-orbit performance. We then summarize calibration challenges encountered during its first 18 months in orbit and the methods used to mitigate their impact on the calibrated radiance spectra distributed to the science community.

## Copyright statement

The author's copyright for this publication is transferred to the California Institute of Technology.

## 1 Introduction

The Orbiting Carbon Observatory-2 (OCO-2) is the first NASA satellite designed to measure the column-averaged carbon dioxide ( $\text{CO}_2$ ) dry-air mole fraction,  $\text{XCO}_2$ , with the accuracy, resolution, and coverage needed to identify and characterize  $\text{CO}_2$  sources and sinks on regional scales over the globe. Surface-weighted  $\text{XCO}_2$  estimates can be retrieved from high-resolution spectroscopic observations of reflected sunlight in near-infrared  $\text{CO}_2$  and molecular oxygen ( $\text{O}_2$ ) bands (Rayner and O'Brien, 2001; Crisp et al., 2004, 2008, 2015; O'Dell et al., 2012). This is a particularly challenging space-based remote-sensing measurement because even the largest  $\text{CO}_2$  sources and sinks produce changes in the background  $\text{XCO}_2$  distribution that rarely exceed 0.25 % (e.g., 1 part per million (ppm) out of the ambient 400 ppm background, <http://cdiac.ornl.gov/>) on spatial scales ranging from that of a satellite footprint (a few square kilometers) to continental scales (Miller et al., 2007). These small variations in  $\text{XCO}_2$  typically introduce changes in the line core-to-continuum intensity ratio for  $\text{CO}_2$  absorption lines that are smaller than 0.1 % in reflected solar spectra, even at spectral resolving powers as high as  $\lambda/\Delta\lambda = 20\,000$ .

To record these small changes in the reflected solar spectrum, OCO-2 carries and points a single instrument that incorporates three imaging grating spectrometer channels that collect high-resolution spectra of reflected sunlight in the 0.765  $\mu\text{m}$  ( $\mu\text{m}$ )  $\text{O}_2$  A-band and in the 1.61 and 2.06  $\mu\text{m}$   $\text{CO}_2$  bands with an unprecedented combination spatial resolution ( $< 3 \text{ km}^2/\text{sounding}$ ), spectral resolving power ( $\lambda/\Delta\lambda > 17\,000$ ), sensitivity (continuum signal-to-noise ratio typically  $> 400$ ), and dynamic range ( $\sim 10^4$ ). Each spectrometer channel collects 24 spectra  $\text{s}^{-1}$  across a narrow ( $< 10 \text{ km}$ ) swath as the observatory flies over the sunlit hemisphere. Coincident measurements from the three spectral channels are combined into "soundings" that are analyzed with a state-of-the-art retrieval algorithm to yield estimates of  $\text{XCO}_2$  and other geophysical quantities (Bösch et al., 2006, 2011; O'Dell et al., 2012; Crisp et al., 2012; Eldering et al., 2016).

OCO-2 was successfully launched from Vandenberg Air Force Base in California on 2 July 2014. After completing a series of spacecraft check-out activities and orbit-raising maneuvers, it was inserted at the front of the 705 km Afternoon Constellation (A-Train; L'Ecuyer and Jiang, 2010) on 3 August 2014. The optical bench and focal planes of the three-channel imaging grating spectrometer were then cooled to their operating temperatures (near  $-6.4$  and  $-152.4^\circ\text{C}$ , respectively), and a series of calibration and validation activities was initiated. Since 6 September 2014, the OCO-

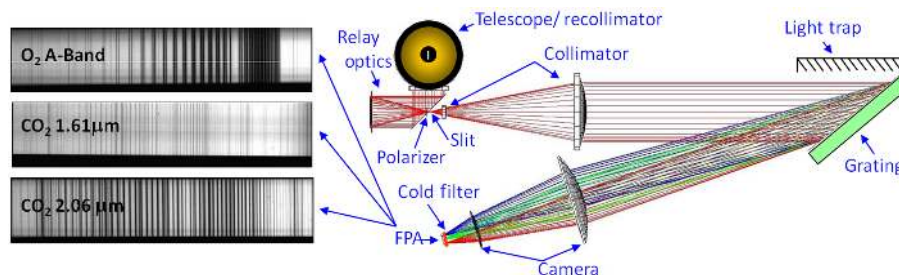
2 instrument has been routinely returning almost 1 million soundings each day over the sunlit hemisphere. These spectra have single-sample continuum signal-to-noise ratios (SNRs) exceeding 400, even over dark ocean surfaces at solar zenith angles as high as  $80^\circ$ . On monthly timescales, 7 to 12 % of these soundings pass the cloud screens and other data quality filters to yield full-column estimates of  $\text{XCO}_2$ . Nadir soundings over land yield  $\text{XCO}_2$  estimates with single-sounding random errors that increase from 0.5 to 1 ppm between the sub-solar latitude and solar zenith angles near  $60^\circ$ . Observations collected near the apparent glint spot, where light is specularly reflected by the Earth's surface, yield  $\text{XCO}_2$  estimates over the ocean with single-sounding random errors near 0.5 ppm at solar zenith angles as high as  $70^\circ$  (Eldering et al., 2016).

Here we describe the instrument performance over the first 18 months of the OCO-2 mission. Sections 2 and 3 provide a brief overview of the instrument design and on-orbit calibration system. Section 4 describes the data products delivered to the science community. Section 5 summarizes the on-orbit radiometric, spectroscopic, and geometric performance of the instrument. Section 6 describes the calibration challenges encountered in orbit and the changes in the calibration and data processing approach implemented to address these issues and to increase the yield and quality of the data products.

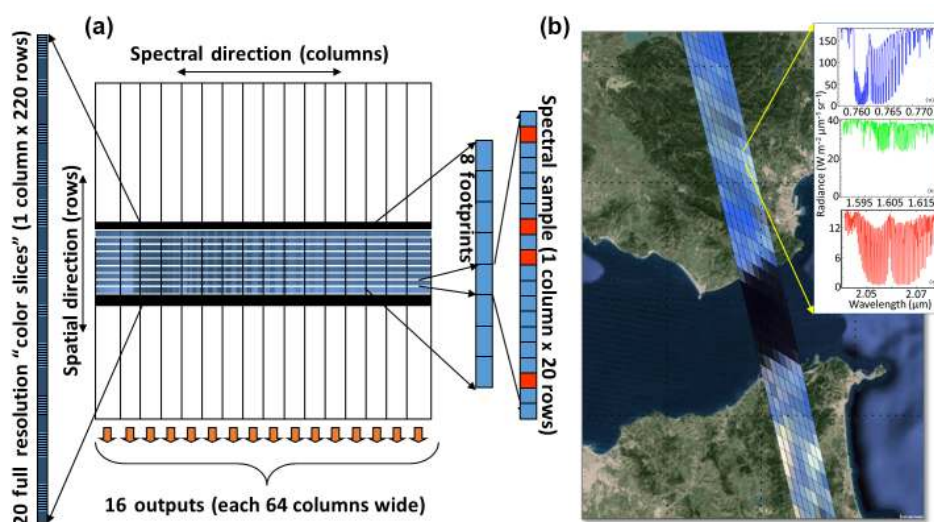
## 2 The OCO-2 instrument

The OCO-2 spacecraft carries and points a single instrument that incorporates three co-boresighted, long-slit imaging grating spectrometers optimized for observing the  $\text{O}_2$  A-band at 0.765  $\mu\text{m}$  and the "weak" and "strong"  $\text{CO}_2$  bands at 1.61 and 2.06  $\mu\text{m}$  (Crisp et al., 2004, 2008; Haring et al., 2004; Pollock et al., 2010). These three channels are designated ABO2, WCO2, and SCO2, respectively. All three spectrometers use similar optical designs and are integrated into a common optical bench assembly (OBA) to improve system rigidity and thermal stability. They share a common, 200 mm focal length, F/1.8 Cassegrain telescope.

The optical path in each spectral channel is shown in Fig. 1, along with examples of recorded spectra. Light entering the telescope is focused at a field stop and then re-collimated before entering a common relay optics assembly. The relay optics use dichroic beam splitters to direct the light to the SCO2, WCO2, and ABO2 channels, in that order. The light then traverses a narrowband pre-disperser filter designed to transmit wavelengths within  $\pm 1\%$  of the central wavelength of the channel and reject the rest (out-of-band transmission  $< 10^{-4}$  of peak transmission). The light is then refocused onto each spectrometer slit by a reverse Newtonian telescope (Haring et al., 2004). Each slit is about 3 mm long and 28  $\mu\text{m}$  wide. These long, narrow slits are aligned to produce co-boresighted fields of view that are about 0.0145 ra-



**Figure 1.** The OCO-2 instrument incorporates three imaging grating spectrometers with similar optical layouts. The major optical components and optical path for a single channel are shown on the right-hand side. The orientation of the narrow entrance slit is indicated by the thin, vertical white line in the secondary mirror at the center of the telescope/re-collimator aperture. Images of spectra recorded by the focal plane arrays (FPAs) in the three spectral channels are shown on the left. A spatially resolved “image” is formed along the slit, while the spectra are dispersed perpendicular to the slit, with wavelength decreasing from left to right.



**Figure 2.** (Left) The illumination and readout approach used for the OCO-2 FPAs, showing the direction of spectral dispersion from left to right and the spatial direction from bottom to top. The  $\sim 160$  illuminated pixels in the spatial dimension are summed into eight  $\sim 20$ -pixel “footprints”. If one or more of the pixels in a footprint is “bad” (red pixels), it is eliminated from the sum. If one or two contiguous pixels in a column are bad, then they are replaced with the average of their good neighbors. If three or more bad pixels are contiguous, then they are replaced with zeros. This algorithm is only applied in the spatial direction – spectral information is never mixed. One of the 20 full-resolution “color slices” is also shown at the left. (Right) The spatial layout of the eight cross-track footprints for a nadir track over the Strait of Gibraltar. The footprints are parallelograms, rather than squares, due to the combined effects of the rolling readout of the FPAs and the spacecraft motion. Each footprint has a cross-track dimension of  $<1.3$  km and a down-track dimension of  $\sim 2.3$  km at nadir. Co-boresighted spectra from the ABO2 (top right), WCO2 (middle right), and SCO2 (bottom right) are collected in each footprint.

dians ( $0.83^\circ$ ) wide at the focal plane. Linear polarizers are installed in front of each spectrometer slit to pass only that component of light polarized perpendicular to the long axis of the slits (e.g., in the direction of dispersion).

Once the light enters a spectrometer slit, it is collimated by a two-element refractive collimator and then dispersed by a gold-coated, reflective, planar, holographic diffraction grating that works in first order. The dispersed light is then focused by a two-element refractive camera lens assembly onto a two-dimensional focal plane array (FPA). A second, narrowband filter, which is cooled to approximately  $-93^\circ\text{C}$ ,

has been installed just above each FPA to further reduce the out-of-band light at wavelengths  $> 2\%$  away from the central wavelength of the channel. This filter also limits the impact of thermal emission from the optical bench, which would otherwise introduce a source of noise in the  $\text{CO}_2$  channels.

The spectrometer optics produce a spatially resolved, two-dimensional image of a spectrum on a  $1024\text{-by-}1024$ -pixel FPA with  $18\mu\text{m}$  by  $18\mu\text{m}$  pixels (Fig. 2a). All three of the FPAs used by OCO-2 are HAWAII-1RG<sup>TM</sup> detectors that were manufactured by Teledyne Scientific and Imaging, LLC, for the original OCO mission (Pollock et al., 2010).

**Table 1.** Instrument properties.

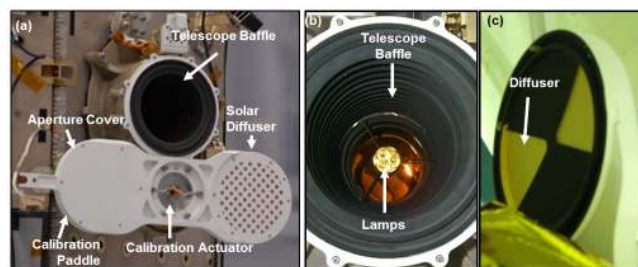
Channel	ABO2	WCO2	SCO2
Spectral rRange ( $\mu\text{m}$ )	0.7576–0.7726	1.5906–1.6218	2.0431–2.0834
Resolving power	17 500–18 500	19 100–20 500	19 700–19 900
Spectral sampling	2.6–3.5	2.3–3.2	2.2–3.2
Dynamic range (photons $\text{m}^{-2} \mu\text{m}^{-1} \text{sr}^{-1} \text{s}^{-1}$ ) (Minimum to maximum measureable signal)	$7.5 \times 10^{16}$ – $7.0 \times 10^{20}$	$2.15 \times 10^{16}$ – $2.45 \times 10^{20}$	$2.15 \times 10^{16}$ – $1.25 \times 10^{20}$

The FPA in the ABO2 channel uses a silicon (HyViSi<sup>TM</sup>) photodetector array, while the CO<sub>2</sub> channels use mercury cadmium telluride (HgCdTe) photodetector arrays with a standard 2.5  $\mu\text{m}$  long-wavelength cutoff. Unlike the FPAs used on the original Orbiting Carbon Observatory (Pollock et al., 2010), the OCO-2 FPAs exhibited negligible residual image (Rosenberg et al., 2016).

The grating disperses the spectrum in the direction perpendicular to the long axis of the slit, illuminating 1016 of the 1024 FPA columns. The four columns at each edge of each FPA are masked and returned as unilluminated reference pixels. The full width at half maximum (FWHM) of the slit image on the FPA (also called the instrument line shape, or ILS) is sampled by 2 to 3.5 pixels in the direction of dispersion. The spectral range, spectral resolving power, number of spectral samples across the ILS FWHM (spectral sampling), and dynamic range for each spectrometer are summarized in Table 1 (Frankenberg et al., 2015; Rosenberg et al., 2016; Lee et al., 2016). The spectral sampling and resolving power vary across the wavelength range of each channel due to the anamorphic magnification introduced by the grating (Pollock et al., 2010).

The ABO2, WCO2, and SCO2 FPAs are cooled to operating temperatures near  $-157$ ,  $-153$ , and  $-152$   $^{\circ}\text{C}$ , respectively, by a pulse tube cryocooler that is connected to the FPAs by thermal straps. The cryocooler rejects its heat through heat pipes that are connected to an external radiator. The temperature of the OBA is maintained near  $-6.4$   $^{\circ}\text{C}$  by a thermal shroud (Haring et al., 2008). The shroud rejects its heat through heat pipes that are connected to a passive external radiator. Because the FPAs and OBA are at very different temperatures, they are linked by a low-conductivity, cryogenic subsystem (Lamborn, 2008; Na-Nakornpanom et al., 2015).

The slits illuminate only the central  $\sim 190$  of the 1024 pixels on each FPA (Fig. 2a). For routine science operations, a 220 (spatial rows) by 1016 (spectral columns) pixel window on each FPA is continuously scanned using a “rolling read-out” method for recording and resetting each pixel on the FPAs to their bias levels (Haring et al., 2004). This readout approach precludes the need for a physical shutter and eliminates spatial gaps between the exposures. The  $0.83^{\circ}$  field of



**Figure 3.** The onboard calibrator (OBC) is attached to the telescope baffle assembly. (a) The OBC is shown in the “open” position for taking science data, with the aperture cover to the left and the solar diffuser to the right. (b) The telescope baffle assembly is shown along with the incandescent lamps used to illuminate the diffuser on the back of the aperture cover. (c) The inner surface of the aperture cover is partially coated with a roughened gold surface that is illuminated by the lamps to produce a spatially uniform illumination field.

view illuminated by the 3 mm long spectrometer slits is sampled by about 190 pixels in the dimension orthogonal to the direction of dispersion (FPA rows). Science measurements are restricted to the central 160 of these 190 pixels, which are fully illuminated by the slits.

The active 220-by-1016-pixel region on each FPA can be read out pixel by pixel, but this readout mode is not practical for routine science observations because it takes 9.333 s for the instrument signal chain to process this many data. The response of each pixel in the 220-by-1016 active area is therefore only recorded during FPA characterization and spectrometer calibration, and not for routine science data collection.

To reduce the FPA readout time and downlink data volume, the 160 rows that are fully illuminated by the slit are divided into eight discrete “footprints”, each consisting of 19 to 21 adjacent spatial pixels along an FPA column (the “spatial direction” in Fig. 2a). The pixels within each footprint are summed together on board by the instrument processor. In this “summed mode”, the active area of the FPAs can be read out at 3 Hz. The eight-footprint “image” of the spectrum that is recorded for each 0.333 s exposure is called a “frame”. The  $\sim 20$ -pixel sum corresponding to a given col-

umn on the FPA is defined as a “spectral sample”. The along-slit angular field of view of each spatially averaged spectral sample is about 1.8 milliradians (mrad), or  $\sim 0.1^\circ$ . The spatial field of view recorded by the 1016 spectral samples in a science spectrum is defined as a “summed footprint”. In addition to reducing the data volume and readout time, this on-board summing reduces the impact of single-pixel read noise (which decreases with the square root of the number of pixels included in the sum), increasing spectral sample SNR in low-illumination conditions.

The angular width of the narrow dimension of the slit is only 0.14 mrad ( $\sim 0.1$  km at nadir), but the focus of the entrance telescope was intentionally built with spherical aberrations to increase the effective full width at half maximum of each slit to  $\sim 0.5$  mrad. This blurring was introduced to simplify the boresight alignment of the three spectrometer slits and to reduce the impact of bad pixels incorporated into spectral samples on the FPAs. Over the 0.333 s integration time for each frame, the spacecraft moves about 2.25 km down-track, yielding eight summed footprints with dimensions of  $\sim 2.25$  km by 1.3 km at nadir when the instrument is operated in “push-broom” fashion with the slit oriented orthogonal to the ground track. Because the rolling readout scans across the 220 active rows sequentially over the 0.333 s exposure while the spacecraft is moving, the surface footprints are shaped like parallelograms rather than squares, even when operated in push-broom fashion (Fig. 2b).

In addition to the eight spatially summed, 1016-element spectra in each frame, each spectrometer returns up to 20 FPA columns (colors) from each FPA without any on-board spatial summing. These “color slices” sample the full, along-slit spatial dimension at single-pixel resolution. Each color slice records a 220-pixel-wide region of the FPA, which includes the full field illuminated by the slit ( $\sim 190$  pixels) as well as a few pixels beyond the ends of the slit that are not directly illuminated. These color slices are used to detect spatial variability within each of the spatially summed spectral samples. Color slice pixels outside the 190-pixel-wide region directly illuminated by the slit can also be used to quantify the thermal radiation and scattered light within the instrument. The FPA columns (wavelengths) for the 20 color slices are specified by commands from the ground.

### 3 The onboard calibration system

An onboard calibrator (OBC) has been integrated into the telescope baffle assembly, which is mounted on an external panel of the spacecraft bus (Fig. 3). This system consists of a calibration “paddle” that carries an aperture cover (lens cap) on one end and a transmission diffuser on the other (Haring et al., 2008). The cover is placed over the telescope aperture to protect the instrument from external contamination during launch and orbit maintenance activities. It is also closed to acquire “dark frames” that are used to monitor the zero-level

offset of the FPAs. The backside of the cover has a diffusively reflecting gold surface that can be illuminated by one of three tungsten halogen lamps installed in the baffle assembly. Lamp 1 is used for routine calibration (every orbit not used for downlink), while lamps 2 and 3 are used less frequently (monthly) to track degradation in lamp 1. These lamp “flat-field” images are used to monitor changes in the relative gain of the individual pixels and spectral samples.

The calibration paddle is rotated  $180^\circ$  from the closed position to place the transmission diffuser in front of the telescope aperture to acquire observations of the Sun. The 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 not aligned. The patterned inner surfaces of the two plates are coated by textured gold surfaces. Sunlight that enters the pin holes on the target (Sun)-facing side of the diffuser is reflected multiple times by the roughened 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 of a moderately dark ( $< 5\%$ ) reflecting Lambertian surface when the telescope is pointed at the Sun.

The 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 absolute radiometric response of the instrument (and solar calibrator). The solar diffuser is 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 ( $\pm 7 \text{ km s}^{-1}$ ) and associated Doppler shifts observed over the illuminated hemisphere. The Doppler-shifted spectra collected over each solar Doppler orbit can be combined to produce a massively over-sampled solar spectrum. These data are used to monitor variations in the ILS and have also been used to refine our understanding of the top-of-atmosphere solar spectrum for the wavelength ranges sampled by OCO-2. The calibration paddle is rotated  $90^\circ$  from either the closed or diffuser positions to open the telescope aperture for routine science or lunar calibration observations. Lunar calibration observations were originally planned to monitor the instrument’s geometric calibration. They are also being used to monitor the degradation of the calibration system’s lamps and solar and lamp diffusers.

### 4 OCO-2 radiance data products

The OCO-2 mission produces a “Level 1B” (L1B) product consisting of full orbits or fractions of orbits of calibrated and geolocated spectral radiances from the ABO2, WCO2, and SCO2 channels. These products are screened for optically thick clouds and then analyzed with a full-

physics remote-sensing retrieval algorithm to yield estimates of XCO<sub>2</sub> and other geophysical parameters that constitute the OCO-2 Level 2 (L2) data products. For routine science operations, the spacecraft points the instrument's bore-sight at the local nadir or in the vicinity of the "glint spot", where sunlight is specularly reflected from the Earth's surface. In a typical nadir-observation orbit, the instrument collects data for about 46.5 min, and each spectrometer collects about 8360 frames (66 880 spatially resolved spectra) at solar zenith angles as high as 85°. For a typical glint-observation orbit, the instrument records science data for about 49.3 min, and each spectrometer collects about 8880 frames (71 000 spatially resolved spectra) at solar zenith angles as high as 84°. Each L1B file or "granule" included in this product contains a record for each spectral sample and footprint in each channel, for each frame the instrument collects while viewing Earth during a single spacecraft orbit. All soundings that pass a series of quality criteria (e.g., spectra, geolocation, and housekeeping data available in all three channels, science aperture open) are included in the L1B product.

The OCO-2 L1B products are delivered to the NASA Goddard Earth Science Data and Information Services Center (GES DISC) for distribution and archiving (<http://disc.sci.gsfc.nasa.gov/OCO-2>). The first L1B products were delivered to the GES DISC in December 2014. During the summer of 2015, all L1B and L2 data collected since routine operations began on 6 September 2014 were reprocessed using Version 7 (V7) of the OCO-2 algorithm. Both "forward" and "retrospective" versions of the Version 7 L1B and L2 products are routinely being delivered to the GES DISC. The forward product, designated V7, uses calibration information based on extrapolations of calibration data that were collected in the recent past. These products are usually delivered to the GES DISC within one week of acquisition. These V7 forward products are generally reliable except when the instrument experiences a significant change in its thermal environment that was outside of the training range used to derive the extrapolated calibration data. The retrospective L1B products, designated V7r, use interpolated calibration data and are therefore expected to be more reliable than the forward products. The V7r products are typically delivered to the GES DISC in 1-month blocks, starting about 5–6 weeks after acquisition. The OCO-2 team recommends that the V7r product be used in applications where high accuracy is needed. The V7 and V7r products are described in greater detail in the OCO-2 data product user's guide and in the L1B algorithm theoretical basis document (ATBD), which are posted along with the products at the GES DISC (<http://disc.sci.gsfc.nasa.gov/OCO-2/documentation/oco-2-v7>).

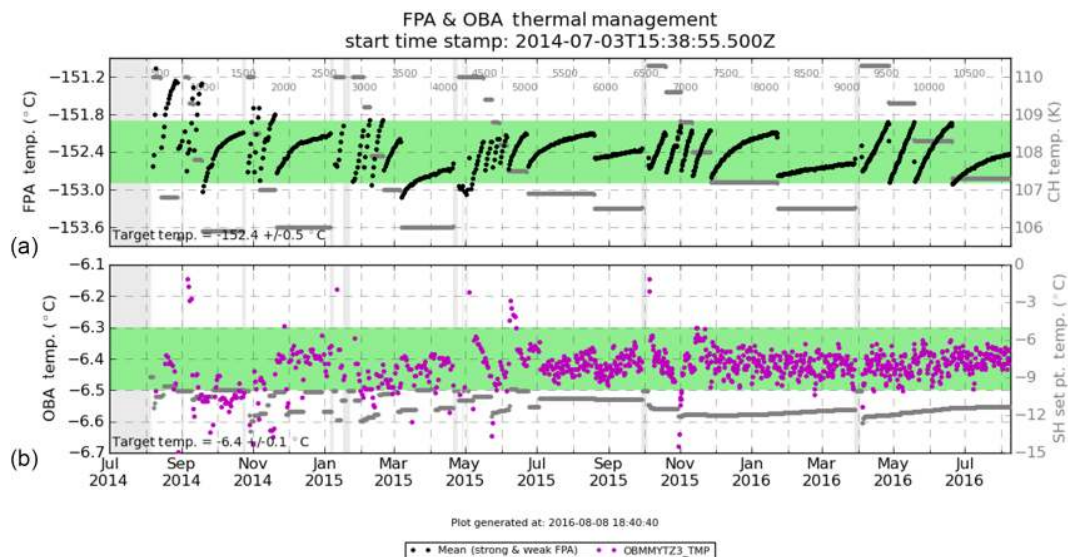
## 5 On-orbit performance

In general, the OCO-2 instrument has performed as expected, yielding spectra of reflected sunlight with an unprecedented combination of spatial resolution, spectral resolution, dynamic range, and SNR. To ensure that the instrument calibration remains stable between calibration opportunities, the OBA temperatures have been controlled to stay within  $\pm 0.2^\circ\text{C}$  of their nominal set points, and the FPA temperatures are controlled to stay within  $\pm 0.4^\circ\text{C}$  of their nominal set points during routine operations (Fig. 4). Instrument decontamination (decon) cycles are performed periodically to remove ice that accumulates on the thermal straps that connect the cryocooler's cold head (CH) to the FPAs. In a typical decon cycle, the OBA and FPAs are warmed to near room temperature (either 12 or 28 °C) for a few days and then cooled back down to their operating temperatures. No science observations can be collected during these  $\sim 1$ -week periods. Fortunately, the frequency of decon cycles has decreased over time as the instrument has lost water to space. Key features of the on-orbit radiometric, spectroscopic, and geometric performance are summarized in the following subsections.

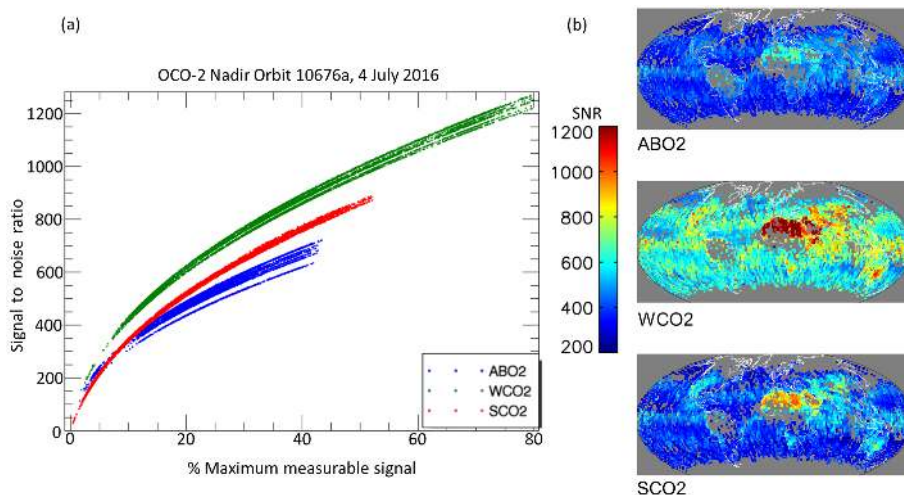
### 5.1 Radiometric performance

The radiometric performance of the OCO-2 spectrometers was established prior to launch through an extensive series of measurements of a calibrated integrating sphere (Eldering et al., 2015; Rosenberg et al., 2016). These and other pre-launch tests (Frankenberg et al., 2015) indicated that the SNR of all three channels substantially exceeded the requirements. The median SNR of individual spectral samples in each channel recorded during a typical nadir orbit are shown in Fig. 5a (see Eldering et al., 2015, and Rosenberg et al., 2016, for a description of the methods used to derive the SNR). In this example, the intensities in the ABO<sub>2</sub>, WCO<sub>2</sub>, and SCO<sub>2</sub> channels spanned  $\sim 40$ , 80, and 52 % of their maximum measurable signals, respectively (i.e., the maximum signal for which the instrument meets its radiometric calibration standards; see Table 1 and Rosenberg et al., 2016). Maps of the continuum SNR for cloud-free soundings used to retrieve XCO<sub>2</sub> estimates in the V7r L2 product for April 2015 are shown in Fig. 5b. Here, the mean continuum SNR values for each channel have been averaged into  $2^\circ$  by  $2^\circ$  bins. With these high SNR values, the range of solar zenith angles or latitudes where XCO<sub>2</sub> can be retrieved is restricted more by clouds and other limitations in the OCO-2 L2 algorithm than by low instrument sensitivity.

The largest SNR values are usually recorded over the Sahara and Australian deserts, but occasionally even brighter values are recorded in glint observations over the ocean at low latitudes. These bright glint scenes sometimes yield continuum intensities substantially exceeding the maximum measurable signal in one or more of the channels, producing



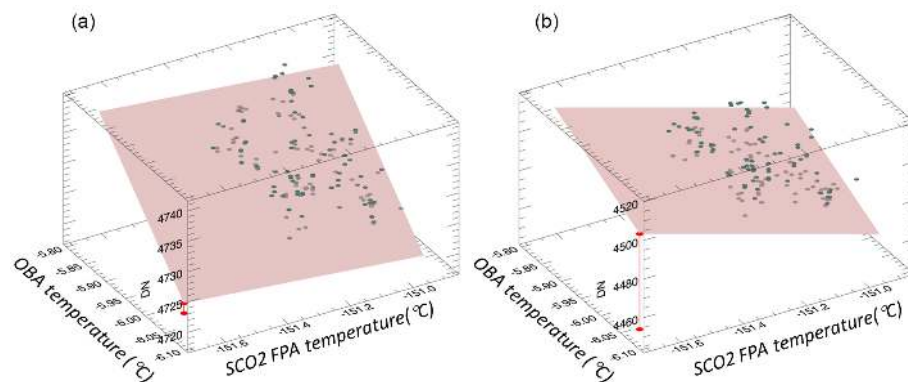
**Figure 4.** Thermal performance of the OCO-2 instrument FPAs (top) and OBA (bottom) are shown as a function of time for normal operations. The FPA temperatures (black dots in top plot) are maintained by controlling the temperature of the cold head (CH) of the cryocooler. The CH temperatures are shown as grey points (horizontal grey lines) in the top plot, and their values are labeled on the right-hand y axis (note: CH temperatures are shifted and stretched relative to the FPA temperatures and expressed in Kelvin rather than degrees Celsius to avoid confusion with the FPA temperatures). During routine operations, the CH temperatures must be adjusted as ice accumulates on the thermal blankets that insulate the thermal straps linking the FPAs to the CH, reducing their efficiency. OBA temperatures (pink dots in bottom plot) are maintained by controlling the temperature of the thermal shroud that encloses the OBA. The set point of the shroud (SH set pt.) is shown as grey points on the bottom plot, and their values are labeled on the right-hand y axis of that plot. The light grey vertical bars indicate decontamination cycles when normal operations are suspended and the FPAs and OBA are raised to near room temperature to remove ice from the thermal straps and FPAs.



**Figure 5.** (a) Median signal-to-noise ratios (SNRs) for individual soundings in the ABO2 (blue), WCO2 (green), and SCO2 (red) channels are shown as a function of the percent maximum measurable signal (see Table 1) for a nadir orbit observed on 4 July 2016. The parallel traces show results for the eight footprints in each channel. A continuum SNR exceeding 200 is typically needed to yield single-sounding random errors near 1 ppm. (b) The mean continuum single-sounding SNR in  $2^{\circ} \times 2^{\circ}$  bins is shown for the ABO2 (top), WCO2 (middle), and SCO2 (bottom) channels for April 2015.

saturated spectra. Saturated channels typically recover within one or two exposures, as expected, with no measurable residual image (see discussion in Rosenberg et al., 2016).

Prior to launch, the absolute radiometric calibration of the instrument was established from observation of an integrating sphere, with reference radiometers validated against Na-



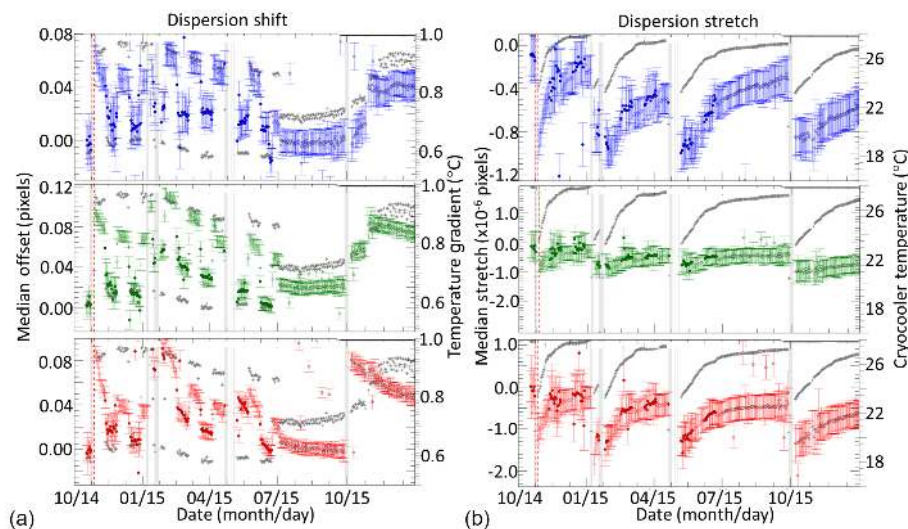
**Figure 6.** The dark coefficient fitting process is illustrated for a typical “well-behaved” spectral sample (left) and a “temperature-sensitive” spectral sample (right) in the SCO2 channel. In both cases, individual spots show the sample response (expressed in digital numbers, DN) as a function of the temperature of the FPA and the OBA. The fitting process assumes that the dark response is a linear function of the FPA and OBA temperatures over the training range. The dark response of the well-behaved sample changes by  $4.7 \text{ DN } ^\circ\text{C}^{-1}$  with respect to the FPA and  $89.9 \text{ DN } ^\circ\text{C}^{-1}$  with respect to the OBA. The dark response of the temperature-sensitive sample changes by  $47.2 \text{ DN } ^\circ\text{C}^{-1}$  with respect to the FPA and  $126.3 \text{ DN } ^\circ\text{C}^{-1}$  with respect to the OBA. If either the FPA temperatures or OBA temperatures move outside the range of values used in these fits, the dark response of the temperature-sensitive sample must be extrapolated and is much less predictable than that of the well-behaved sample.

tional Institute of Standards and Technology (NIST). This system was required to meet a 5 % absolute radiometric requirement but yielded much smaller uncertainties (1.6, 3.2, and 2.9 % in the ABO2, WCO2, and SCO2 channels, respectively; see Rosenberg et al., 2016). Since launch, the radiometric performance has been routinely monitored by deploying the solar diffuser in front of the instrument aperture to collect observations of the Sun, or by placing the lamp diffuser in front of the aperture to collect both dark (lamp off) and flat-field (lamp on) exposures. Dark, solar, and lamp diffuser observations are collected in “summed” and or “single-pixel” mode on alternate orbits that are not used for downlink or other maintenance activities. These routine radiometric calibration activities are augmented by observations of the Moon, collected for near full and gibbous phases each month. In addition, observations of a backup calibration lamp are collected at roughly monthly intervals, and targeted observations of well-calibrated surface vicarious calibration sites – such as Railroad Valley, Nevada – are collected periodically when the site is cloud free.

The radiometric calibration of each spectral sample can be described in terms of a dark offset and a gain, both of which were calibrated prior to launch (Rosenberg et al., 2016). The dark offset is defined as the response of a sample when no external illumination is entering the instrument. This component of response must be updated frequently in orbit because the dark offset of a few pixels is sensitive to small (millikelvin, mK) changes in the temperature of the FPA. The dark offset in the WCO2 and SCO2 channels is also sensitive to the temperature and associated black-body emission of the OBA.

To characterize dark offset changes, either summed or single-pixel measurements are collected for 1 to 3 min with the aperture door closed on every orbit where routine calibration measurements are collected. These dark data are used to update the dark offset correction coefficients described by Rosenberg et al. (2016). For the forward V7 L1B product, FPA and OBA temperature trends from the past 10 days (typically 146 orbits) are extrapolated to predict the dark offset corrections for the following week. For the retrospective product, V7r, temperature-dependent dark offsets corrections are derived using measurements acquired during the time period when the data were acquired.

Typically, OBA temperatures vary by around 100 mK each day due to orbit-to-orbit changes in spacecraft operations (Fig. 4). FPA temperatures vary much more slowly ( $< 5 \text{ mK day}^{-1}$ ), but their temperatures increase as ice accumulates on the thermal straps that connect the FPAs to the cryocooler CH. Periodic adjustments as large as 1 K are made to the cryocooler CH temperatures to compensate for these changes. In the forward (V7) product, these large corrections require the dark offset calibration coefficients to be extrapolated far outside the training range. For most spectral samples, this extrapolation has minimal impact because the dark offset is relatively insensitive to temperature, while a few samples have much greater temperature sensitivity (Fig. 6). If their dark offsets are not accurately described by the extrapolated correction coefficients, errors can be introduced in the L1B product that can degrade the spectral fits performed by the L2 algorithm and introduce biases in the retrieved products. CH adjustments and other unanticipated FPA and OBA temperature corrections do not affect the retrospective V7r



**Figure 7.** The time evolution of the dispersion shift/offset (left column) and dispersion stretch (right column) is shown for the ABO2 (top, blue symbols), WCO2 (middle, green symbols), and SCO2 (bottom, red symbols) channels. Both are shown on the left-hand ordinate of each plot. The shifts are expressed in pixels, while the stretch is expressed as a ratio with the fixed dispersion stretch term, subtracted from 1, in parts per million. The shifts are well correlated with temperature gradients across the OBA (dark grey crosses and right-hand ordinate), which can introduce grading tilts and other distortions in the OBA. The stretch term is correlated with the cryocooler temperature (dark grey symbols and right-hand ordinate). These shifts appear to be correlated with time between decon cycles, indicated as vertical grey bars.

products because they use interpolated rather than extrapolated data to derive dark offset correction coefficients.

In addition to dark offset variations, routine on-orbit calibration measurements have revealed changes in the radiometric performance due to increased numbers of bad pixels in the WCO2 and SCO2 FPAs, ice accumulation on the ABO2 FPA, and much smaller changes thought to be due to optical coating degradation. These phenomena and the methods used to characterize and mitigate their impacts are described in Sect. 6.1 and 6.5, respectively. To date, besides the effects of bad pixels, on-orbit measurements provide no compelling evidence for changes in the overall sensitivity or gain linearity of the FPAs.

## 5.2 Spectroscopic performance

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., 2016). The dispersion relation for each spectrometer specifies the mapping of the wavelengths of light onto specific pixels on the FPAs. The dispersion relation for each footprint in each spectral channel was derived prior to launch (Lee et al., 2016) and is specified in terms of a fifth-order polynomial, the coefficients of which are listed in the L1B data files. 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. Changes in the focus of spectrometer collimator or camera assemblies could also produce changes in both the dispersion and the ILS.

The spectroscopic performance has been routinely monitored in orbit using measurements of the positions and shapes of solar lines observed through the solar diffuser. Additional insight has been gained by monitoring changes in the fits to atmospheric spectra generated by the L2 retrieval process. These tests reveal small variations in the dispersion associated with changes in the instrument thermal environment and with ice accumulation on the thermal straps that connect the cryocooler cold head to the FPAs. 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.

Data collected during the first 18 months of the OCO-2 mission indicate small dispersion changes that can be compensated by changes in a wavelength offset or “shift” and wavelength “stretch” coefficients in the dispersion polynomial. The time evolution of the shift and stretch coefficients is shown in Fig. 7. Prior to 2 July 2015, when nadir and glint observations were collected continuously for full, 16-day ground track repeat cycles, the instrument thermal environment equilibrated to distinct “nadir” and “glint” states that are clearly seen in the dispersion shift and stretch coefficients. Between June and November, when the observing mode switched between glint and nadir observations on alternating orbits, the instrument thermal environment remained more stable over time. On 12 November 2015, when the ob-

serving strategy was further optimized to always collect glint observations over ocean-only orbits, a new, stable thermal equilibrium was established.

In addition to its dependence on observing mode, Fig. 7 shows that the dispersion also changes with time between instrument decon cycles. These small changes were not characterized prior to launch but are thought to be driven by the following mechanism. As ice accumulates on the thermal straps, thermal losses between the CH and FPAs increase and the cryocooler duty cycle increases to compensate. The accompanying increase in thermal load creates a thermal gradient across the cryocooler and the OBA structure, which increases with ice accumulation. This gradient introduces stresses on the thermal straps that connect the CH to the FPAs. These stresses introduce small ( $< 1/10$  of the width of a pixel) lateral displacements (shifts) in the ABO2, WCO2, and SCO2 channels, with amplitudes that appear to depend on the distance between the FPA and the cryocooler CH. They also appear to produce small tilts in the FPAs, relative to the optical axis of the system, introducing small changes in the dispersion stretch.

In the V7 product, the small changes in the dispersion shift and stretch coefficients are derived and corrected as part of the spectral fitting process in the L2 retrieval algorithm. These changes are therefore not reported in the L1B V7 or V7r products delivered to the GES DISC. Using data collected over the past 18 months, the OCO-2 calibration team has identified robust relationships between dispersion shift and stretch coefficients and CH and OBA temperatures that are routinely monitored. These insights are being used to track dispersion changes so that they can be included in future versions of the L1B product. This should simplify the use of the L1B product by future users. The forward L1B product will include predicted dispersion shift and stretch values that are based on the instrument thermal environment. The retrospective L1B product will include refined estimates of the dispersion coefficients retrieved by the L2 algorithm.

### 5.3 Geometric performance

To retrieve XCO<sub>2</sub> estimates with accuracies near 0.25 %, the optical path length must be known to a comparable accuracy. An accurate geolocation of each sounding footprint is essential for defining the optical path length associated with each spectrum. 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, and the relative pointing of the fields of view of the eight footprints in the three spectrometer channels must be known. The location of the observatory along its orbit track is determined by a Global Positioning System (GPS) sensor. The spacecraft orientation relative to fixed stars is determined by a star tracker. 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. In addition, a series of on-orbit calibration activities were implemented to verify the relative alignment of the spectrometer boresights, to refine our knowledge of their alignment with the star tracker, and to provide an end-to-end verification of the geolocation software.

The relative alignment of the star tracker and the instrument boresight and the co-alignment of the instrument slits were tested with observations of the Moon. In these tests, the nominal center of the instrument field of view was pointed at an inertial point in space (a distant star field) that was expected to pass behind the gibbous lunar disk as the satellite traversed the nightside of the orbit. These lunar tests verified that the instrument slits had not moved relative to each other by a measurable amount ( $\sim 0.07$  mrad, or 50 m at nadir) between the pre-launch testing and orbit. They also indicated that the alignment of the star tracker and instrument was within specification ( $< 1^\circ$ ) and provided correction factors that were incorporated into the geometric calibration algorithm. The second step in the validation of the on-orbit geometric calibration involved observations of shoreline crossings for routine nadir and glint observations. This effort yielded geolocation errors no larger than 0.25 mrad (0.2 km at nadir), which is much smaller than the specification ( $0.9 \text{ mrad axis}^{-1}$ ,  $3\sigma$  or  $\sim 0.9$  km worst case at nadir).

## 6 On-orbit instrument calibration challenges

To fully exploit the information content of the high SNR OCO-2 spectra, they must be calibrated to a high degree of accuracy. For example, the relative radiometric performance (defined in terms of the dark offset, gain, and gain linearity; see Rosenberg et al., 2016) of spectral samples within a given channel must be known to within 0.1 % of the continuum brightness across the spectral range of each channel to fully exploit the spectrally dependent information in each sounding. The spectral calibration requirements (dispersion, ILS width, ILS shape) are equally demanding, since uncertainties in these parameters lead to biases in retrievals (Connor et al., 2008, 2016). While the OCO-2 instrument was characterized and calibrated prior to launch (Frankenberg et al., 2015; Rosenberg et al., 2016; Lee et al., 2016), a number of calibration challenges were encountered during the first few months of operation. These included increased numbers of bad pixels on the CO<sub>2</sub> FPAs, transient artifacts introduced by cosmic ray hits, spectral discontinuities for spatially inhomogeneous footprints, and time-dependent changes in the sensitivity of the ABO2 channel. The methods used to identify, characterize, and mitigate the impacts of each of these issues are described in the following sections.

### 6.1 Bad pixels and bad samples

The FPAs used by OCO-2 are flight-spare units from the original OCO mission, which were delivered in 2006. A few

percent of the 162 560 pixels in the 1016-by-160-pixel active area that is used to record science data either are dead or respond to light or temperature 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 on board or they will contaminate the 20-pixel spectral samples that constitute the spectra returned for each of the eight footprints and three spectrometer channels. 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 samples” in the L1B product (see variable L1BSc/InstrumentHeader/snr\_coef) and should only be used with caution (or ignored) in higher-level processing.

The bad-pixel identification process was started during pre-launch testing but has to be updated periodically as the FPAs age. Many additional bad pixels formed during the 2.5-year instrument storage period, which was imposed by the need to replace the OCO-2 launch vehicle after the loss of the Glory mission. The performance of other pixels has degraded in orbit due to cosmic radiation, thermal cycling of the FPAs, and other factors. The HgCdTe FPAs used in the WCO2 and SCO2 channels accumulated a large number of new bad pixels, particularly on the long-wavelength ends of their spectral ranges while the instrument was in storage prior to launch (Fig. 8). More than 2.5 % of the pixels in the 160-by-1016-pixel area used for recording science data are now classified as bad on these two FPAs (Table 2). In contrast, the ABO2 FPA developed very few new bad pixels while in storage. The rate of bad-pixel formation in all three FPAs has been much slower since the observatory was launched and the FPAs were cooled to their operating temperatures. This experience is consistent with the experience of the James Webb Space Telescope (JWST) Detector Degradation Failure Review Board (Rauscher et al., 2012), which documented rapid degradation of this class of HgCdTe FPAs at room temperature and much slower degradation at cryogenic temperatures.

To identify bad or degraded pixels or spectral samples, the calibration team routinely collects dark frames and lamp flat-field frames using single-pixel mode, which returns all 225 280 pixels in the 220-by-1024 active region of each FPA. These single-pixel calibration data are analyzed to identify bad pixels, which are added to a “bad-pixel map” for each FPA. This map is a 1024-by-220 byte array, where a value of 0 indicates that a pixel is good and a value of 1 indicates that a pixel’s signal should not be included in a spectral sample. Revised bad-pixel maps are periodically uploaded to the spacecraft for use in the pixel-summing process.

As noted above, the spectral samples are created on board by summing 19 to 21 adjacent pixels in a given FPA column. The algorithm is quite simple and does not apply bias or gain corrections to individual pixels prior to incorporating them into a sample because the instrument controller was not fast enough to perform this calculation on board. It does,

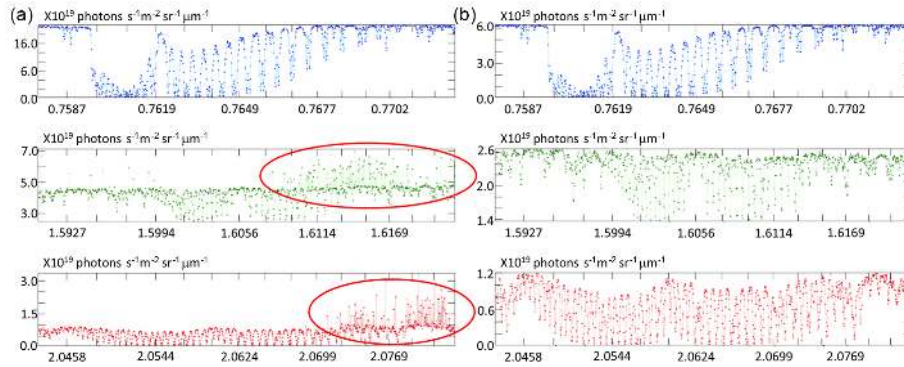
however, assign different weights to pixels, based on their proximity to bad pixels. A single or double bad pixel is replaced by the average of the closest adjacent normal pixels. If three or more adjacent pixels are bad within a single spectral sample, those values are excluded from the sum and the weights of the other pixels are renormalized. This is a good approximation because the spatial imaging of the telescope was intentionally blurred to yield a point spread function with a FWHM of three to six pixels. In this algorithm, a pixel within a given sample can therefore have a weight of 0 (bad), 1 (normal, with no adjacent bad pixels), 1.5 (one adjacent bad pixel), 2 (2 adjacent bad pixels), 2.5, or 3.

Early in the OCO-2 mission, the identification of bad pixels was a high priority because large numbers of bad pixels appeared on the CO<sub>2</sub> FPAs during the pre-launch storage period. These bad pixels produced artifacts in the spectra that compromised their utility. As the bad-pixel maps improved, a larger fraction of the spectral samples in each spectrum could be used to generate higher-level products. Because the ABO2 FPA did not degrade significantly during the pre-launch storage period, its bad-pixel map has not been updated since launch. In contrast, the bad-pixel maps for the two CO<sub>2</sub> FPAs were updated four times between September 2014 and March 2015 as the bad-pixel population was better characterized (Table 2). The bad-pixel map version used for each granule is specified in the metadata of each file in the variable “BadPixelMapVersionNum”<sup>1</sup>.

In addition to the bad-pixel map updates, the calibration team also maintains a bad-sample list, which is applied later in data processing to mask additional spectral samples with inadequately corrected bad pixels. This approach provides more flexibility than the bad-pixel maps alone, because it does not permanently mask spectral samples whose performance is better calibrated through further analysis (Note that, while bad samples can be recovered through further calibration, individual pixels labeled as “bad” are not subsequently relabeled as “good”). The bad-sample lists are updated with each bad-pixel map update, or when a sample’s performance is altered by a cosmic ray event, decon thermal cycle, or other factors. The bad-sample list is also used to mask spectral samples at the edges of the FPA where the ILS could not be adequately characterized and samples at the short-wavelength end of each band that are compromised by optical distortion. These latter two factors account for the majority of the samples included in the bad-sample list.

The most significant change to the bad-sample list was made on 10 November 2014 (Table 3), when 565 and 517 samples were removed from the list in the WCO2 and SCO2 channels, respectively, based on further characterization of in-flight data. The only change to the ABO2 channel was the addition of five samples to the list on 13 February 2015. In addition, the WCO2 bad-sample list was increased by one

<sup>1</sup>Note: the bad-pixel map is applied only to summed spectral samples. It is not applied to the full-resolution color slices.



**Figure 8.** Spectra recorded in footprint 3 of the ABO2 (top), WCO2 (middle), and SCO2 (bottom) channels on 8 August (left) and 10 September (right) show the impact of an improved dark calibration. The 8 August spectra were calibrated using FPA dark response recorded during the pre-launch thermo-vacuum tests at JPL in 2012. The 10 September spectra were calibrated using dark data recorded in flight. The revised dark data and bad-pixel maps dramatically improve the quality of the spectra at longer wavelengths in the WCO2 and SCO2 channels.

**Table 2.** Bad-pixel map version and number of bad pixels. Date format: month/day/year.

Start date	Bad-pixel map version ABO2/WCO2/SCO2	Start orbit	ABO2 # bad pixels	WCO2 # bad pixels	SCO2 # bad pixels
3/30/2012	5/5/5	pre-flight	853	1606	1400
9/5/2014	5/6/6	955	853	2725	2164
11/10/2014	5/8/8	1913	853	3631	2859
11/21/2014	5/9/9	2074	853	4371	3699
2/13/2015	5/10/10	3297	853	4520	4414

sample on 1 July 2015, and the SCO2 bad-sample list was increased by a single sample each time on 13 October 2015 and 10 November 2015 to respond to new bad samples caused by cosmic ray events. The WCO2 and SCO2 channels were changed significantly (adding 49 and 155 samples to the bad-sample list, respectively) on 1 January 2016 in response to improved sample dark correction. On 1 March 2016, the bad-sample list was updated to add one more sample to the list in the SCO2 channel. A complete list of bad-sample changes is provided in Table 3, where we list the number of samples rejected, out of the 8128 available in each channel. The distribution of bad samples is shown as a function of wavelength (L1B column), footprint, and channel in Fig. 9.

With each new bad-pixel map, new bad-sample lists and calibration tables were constructed. Each new bad pixel required modifications to the radiometric gain of its spectral sample. For V7 L1B products, the gain and offset adjustments were tested by performing L2 retrievals on a large test data set to identify persistent outliers. Spectral samples that continued to produce large spectral residuals were added to the bad-sample list. This method for identifying new bad samples has been refined over the first 18 months of the mission and is now a routine part of the L1B validation process (Table 3).

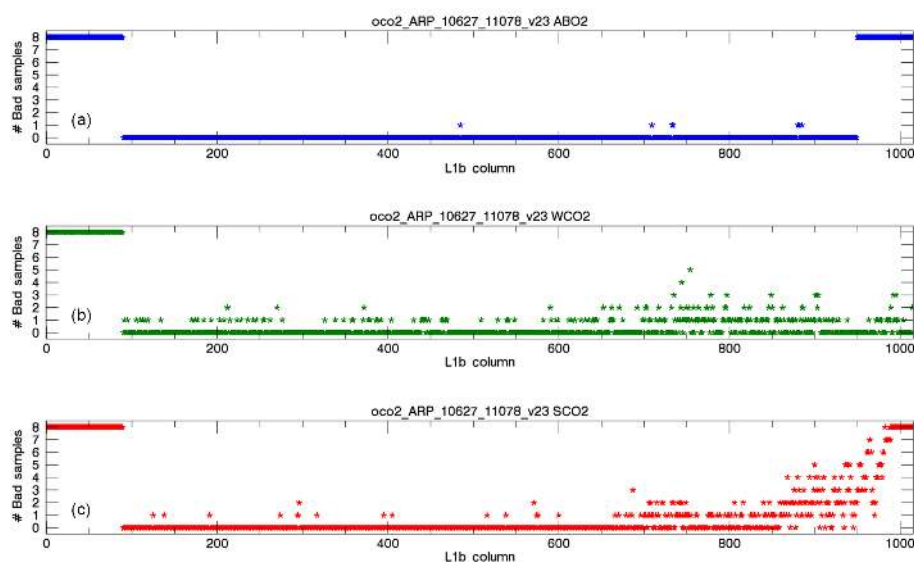
## 6.2 Transient cosmic ray artifacts

Cosmic rays rarely produce permanent damage to the OCO-2 FPAs, but energetic particles with a range of energies produce ion trails as they traverse the FPAs. Because the electrons in these ion trails are indistinguishable from those produced by photons, they produce spurious intensity spikes in individual exposures. The largest cosmic ray spikes are seen in the ABO2 channel, because the silicon photodiode array used in the ABO2 FPA is physically thicker ( $\sim 100\mu\text{m}$ ; Bai et al., 2008) and more sensitive to cosmic rays than the thinner ( $5\text{--}10\mu\text{m}$ ; Beletic et al., 2008) cadmium–zinc–tellurium (CdZnTe) substrates used for the FPAs in the CO<sub>2</sub> channels.

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 the base of the inner Van Allen radiation belt extends to altitudes below the 705 km orbit altitude of OCO-2 (Stassinopoulos and Raymond, 1988). The approximate geographical extent of the SAA impacts on the OCO-2 sensors can be seen in Fig. 10. Over this region, up to 2 % of the 1024 spectral samples in an O<sub>2</sub> A-band spectrum can be contaminated in a single 0.333 s exposure. An example of cosmic ray artifacts in a spectrum is shown in Fig. 11, where large cosmic ray spikes are seen above the A-band continuum across the entire spectrum. All but the

**Table 3.** Bad-sample list change history.

Start orbit	ABO2 samples	WCO2 samples	SCO2 samples	Reason	Date
958	2498	2241	2471	Bad-pixel map 6	2014-09-06
1913	2498	1676	1954	Bad-pixel map 8	2014-11-10
2074	2498	1695	1934	Bad-pixel map 9	2014-11-21
3297	2503	1676	2196	Bad-pixel map 10	2015-02-13
5297	2503	1677	2196	Cosmic ray event	2015-07-01
6820	2503	1677	2197	Cosmic ray event	2015-10-13
7227	2503	1677	2198	Cosmic ray event	2015-11-10
7977	2503	1726	2353	Sample dark correction	2016-01-01
8851	2503	1726	2354	Single new bad sample	2016-03-01



**Figure 9.** The bad-sample distribution is shown as a function of the L1B column number and footprint for the ABO2 (top), WCO2 (middle), and SCO2 (bottom) channels. Bad samples in footprints 1 through 8 are indicated by a \* in the corresponding y axis value. All columns with all good samples are indicated by a star in y axis value, 0. All samples in columns less than  $\sim 180$  have been marked as bad on all three FPAs due to optical aberrations that are not well characterized on that end of the array. Samples near the right-hand side of the FPAs have been marked as bad because the ILS is not well characterized there. A larger number of bad samples are marked on the right-hand side of the WCO2 and SCO2 FPAs because more bad pixels have appeared on that side of these FPAs.

brightest of these artifacts are eliminated when the FPA is reset for the next exposure.

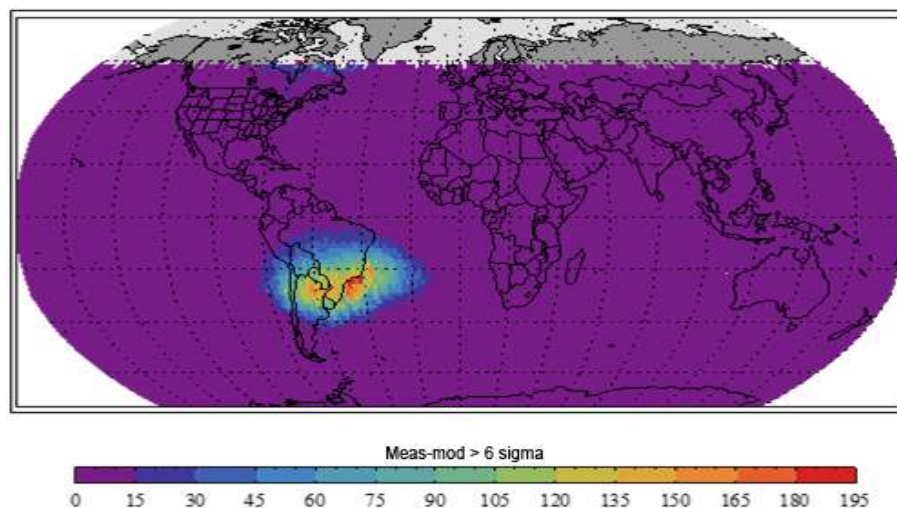
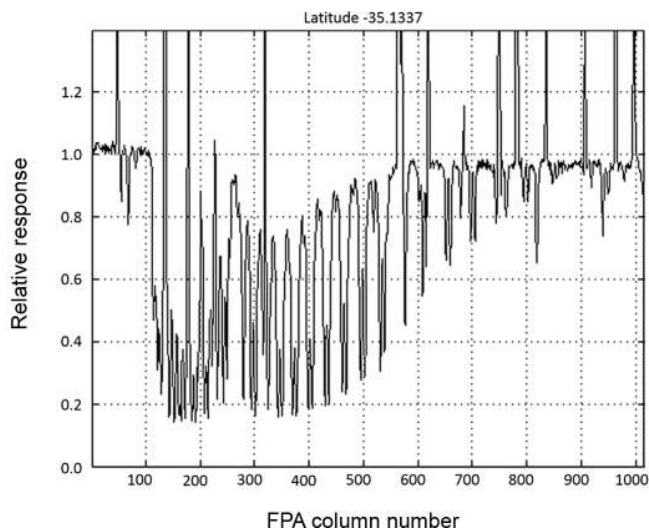
An algorithm was developed to identify and screen cosmic ray artifacts in the L1B products delivered to the science community (Eldering et al., 2015). For each orbit granule, this algorithm performs a singular value decomposition of all spectra collected in each footprint in each band. A least-squares approach is then used to fit the 40 leading eigenvectors to each spectrum and then identify the spectral samples that are outliers from the fit. These outliers (positive spectral residuals) are flagged, based on the number of standard deviations between the fit and the recorded spectrum for each

sample. The degree of disagreement is recorded in the L1B file, in the group “SpikeEOF”.

The OCO-2 L2 retrieval algorithm uses the weighted residuals in the L1B file to flag individual spectral samples as “contaminated”. For the V7 product, samples are assumed to be contaminated if their weighted spectral residual exceeds  $+6\sigma$ , and if the sounding falls within a geographic region around the SAA ( $-50$  to  $0^\circ$  latitude,  $-90$  to  $10^\circ$  longitude). A one-sided threshold value is used because cosmic ray hits can only cause a positive anomaly in measured spectra. The V7 L1B ATBD (Eldering et al., 2015) advises users to adopt a similar approach for flagging spectral samples contaminated by cosmic rays and to apply this threshold only in the

**Table 4.** Clocking adjustment columns in each spectral channel.

Channel	ABO2	WCO2	SCO2
Clocking adjustment columns	123, 363, 603, 843	51, 475, 899	87, 211, 335, 459, 583, 707, 831, 955

**Figure 10.** The number of spectral samples in the ABO2 channel that are contaminated by  $10\sigma$  cosmic ray events clearly shows the geographical extent of the South Atlantic Anomaly (SAA) for data collected in November 2014–January 2015.**Figure 11.** A raw ABO2 spectrum showing the impact of a cosmic radiation event over the South Atlantic Anomaly. Cosmic rays produce spurious positive spikes that contaminate one or more adjacent pixels.

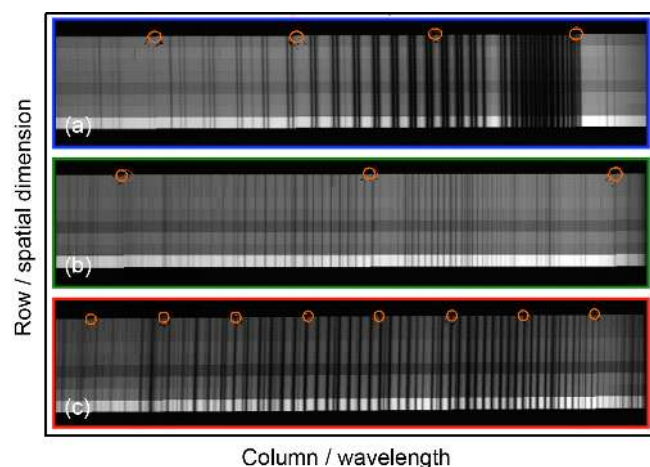
vicinity of the SAA. The V7 cloud screening (Taylor et al., 2016) did not incorporate a cosmic ray flag and thus often failed to converge for many soundings in the SAA region. These soundings were flagged as bad and not processed in this version of the product. To address this issue, a cosmic

ray screening step will be performed to correct or remove corrupted spectral points before running the cloud screening algorithms in future data products.

### 6.3 Radiance discontinuities due to FPA rotation

For imaging grating spectrometers, the spectrometer slits, the grooves on the diffraction gratings, and columns of the FPAs must be well aligned, and other optical distortions (keystone, smile) must be minimized to ensure that a fixed series of rows on the FPA will sample the same angular field of view (or spatial footprint) throughout the spectral range recorded by the FPA. For the OCO-2 instrument, perfect alignment (clocking) of the FPAs with the other optical components was precluded by 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 (Fig. 12). The clocking angles are small ( $\sim 0.3^\circ$  counterclockwise,  $0.2^\circ$  clockwise, and  $0.5^\circ$  clockwise, for the ABO2, WCO2, and SCO2, respectively) but significant because a given geographic position does not map onto a single row of pixels on the FPA but instead varies, roughly linearly, with spectral position across the FPA.

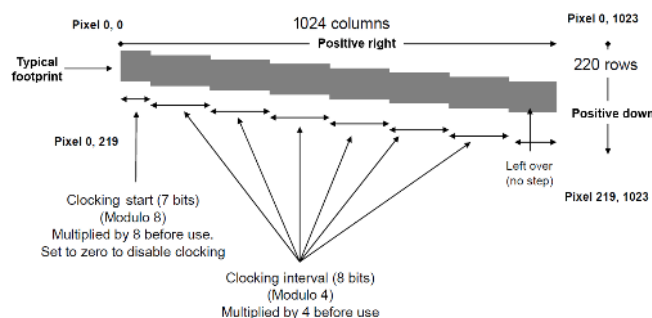
To compensate for the rotation of the OCO-2 FPAs, and record the same spatial information across the entire spectral range, the starting row index for each of the 1016 spectral samples corresponding to a given footprint can be adjusted in increments of one pixel. The columns where



**Figure 12.** Images of the ABO2 (top), WCO2 (middle), and SCO2 (bottom) channels in “summed mode” taken from the OCO-2 first light frame over Papua New Guinea. The eight spatially summed footprints in each band are shown from bottom (footprint 1) to top (footprint 8). In this 0.333 s frame, footprint 1 was contaminated by a cloud and appears brighter than the rest. This frame clearly shows the rotation, or “clocking”, of each FPA columns with respect to the dark O<sub>2</sub> and CO<sub>2</sub> absorption lines. The O<sub>2</sub> lines are tilted counterclockwise by 0.3°, while the CO<sub>2</sub> lines in the WCO2 and SCO2 channels are tilted clockwise by ~0.2 and 0.5°, respectively. The columns where the one-pixel clocking adjustments are applied are highlighted with red circles.

these clocking adjustments occur are recorded in the L1B data files delivered to the community (variable name: clocking\_shift\_color\_indicator) and listed in Table 4. A single pixel corresponds to about 1/20th of the area of a summed footprint (Fig. 2). Adjusting the footprint boundaries in single-pixel increments works well in spatially homogenous scenes but can produce radiance discontinuities in spectra of scenes with strong intensity variations, especially if large spatial gradients in brightness occur near the edge of a footprint. In such scenes, the radiance discontinuities will typically have opposite signs in adjacent footprints. These clocking discontinuities are corrected as part of the L1B calibration process.

The clocking correction algorithm enforces continuity in radiance across single-pixel footprint shifts to compensate for spectral discontinuities associated with single-pixel adjustments (Fig. 13). This algorithm exploits the spatial information provided by the single-pixel color slices described in Sect. 2. Although the spectra recorded for each spatial footprint consist of 1016, ~20-pixel spectral samples, the color slices record the spatial information at single-pixel resolution across each footprint for up to twenty of the 1016 columns (wavelengths) across each FPA. The clocking correction algorithm uses the single-pixel radiances recorded in the color slices to map spatial variation across the spectral



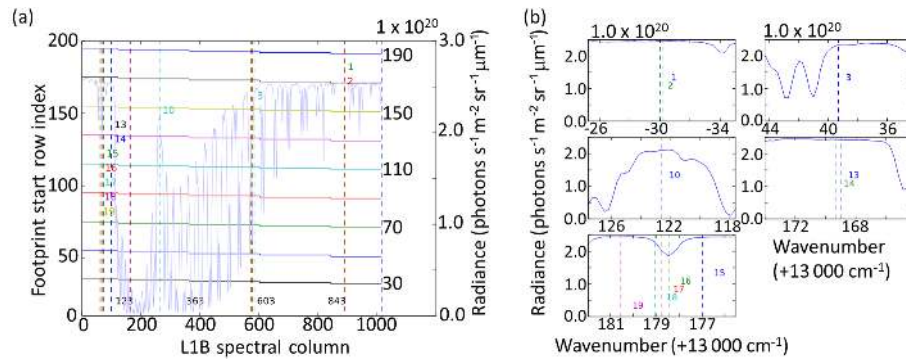
**Figure 13.** The clocking scheme used to sample the same spatial information (vertical dimension) across the entire spectral range (horizontal) for a rotated FPA.

range, particularly near those columns where clocking jumps occur (Fig. 14a).

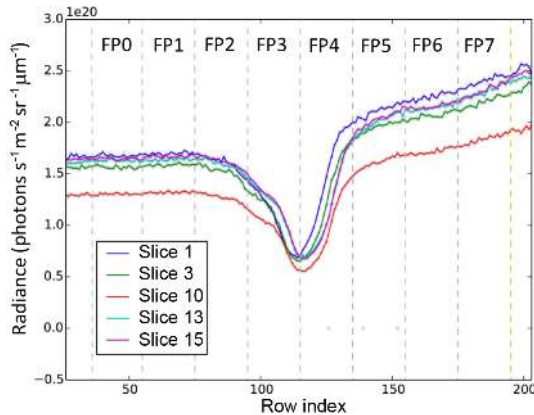
As noted in Sect. 2, the locations (columns) for the color slices on each FPA can be changed by commands from the ground. The most recent update was 17 October 2014 (orbit 1567). The color slice positions are listed in the L1B data files (see variables: ColorSlicePositionO<sub>2</sub>, ColorSlicePositionStrongCO<sub>2</sub>, ColorSlicePositionWeakCO<sub>2</sub>). In principle, placing the color slices adjacent to the columns with the clocking jumps would be ideal, but this is sometimes problematic because some of these columns include sharp atmospheric or solar absorption features, and these features can be shifted by up to three columns by Doppler shifts associated with the spacecraft motion and observing geometry. If a color slice is placed in a spectral region where the radiance changes rapidly (e.g., near an absorption line), spectral variations could be misinterpreted as artifacts associated with spatial inhomogeneity across a footprint (Fig. 14). Therefore, the color slices used by the clocking correction algorithm were located in columns where the spectral variations were small (i.e., in continuum regions) wherever possible.

The availability of continuum regions varies substantially from channel to channel. For the ABO2 FPA, continuum regions were selected just beyond the weaker absorption lines near the edges of the FPA as well as in the continuum near the center of the band. A similar approach was adopted for the WCO2 channel. The deeply saturated SCO2 channel offers no true continuum regions, so the color slices used by the clocking algorithm are located in the most slowly varying regions available near the edges and center of the CO<sub>2</sub> band.

The number of clocking adjustments needed in each channel depends on the rotation angle of the FPA relative to the slit. It therefore varies from channel to channel, ranging from eight discontinuities in the SCO2 channel to only three in the WCO2 channel (Fig. 12). Figure 14 illustrates the choice of color slices for the ABO2 channel. In these figures, only those color slices that are labeled by a column index are used by the clocking algorithm. Other color slices have been lo-



**Figure 14.** (a) Schematic illustrating the clocking for the O<sub>2</sub> A-band along with the current color slice selection. Horizontal, piecewise-constant solid curves delineate footprint boundaries. The discrete jumps in the footprint boundaries produce the clocking discontinuities. Dashed vertical lines indicate color slice selection. Only those vertical lines that are numbered are used in the clocking correction algorithm. Other color slices in strongly absorbing regions (i.e., those near columns 200 and 600) were intended for cloud screening applications. A typical A-Band spectrum is overlaid (blue line) for reference. (b) The color slices used in the clocking correction for the ABO2 band are aggregated into five “groups”. Each panel depicts a single group of color slices (dashed lines) used to identify radiance discontinuities. Spectral features for a “typical” O<sub>2</sub> A-band spectrum are shown as a solid blue line.



**Figure 15.** O<sub>2</sub> A-band radiances sampled 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 dashed vertical lines represent the boundaries of footprints (labeled FP0-FP7). 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 spectral discontinuities in the measurement.

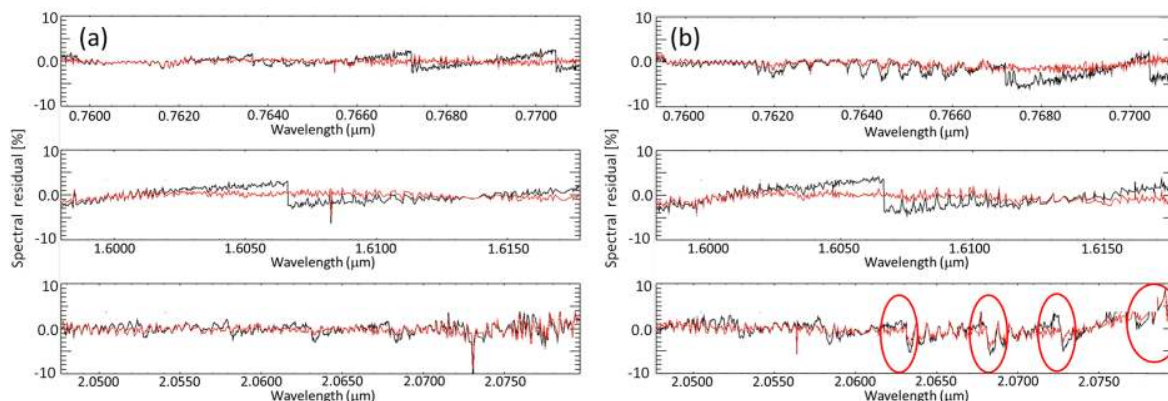
cated within strong absorption features to assess the spatial inhomogeneity due to clouds and aerosols or to monitor solar lines used in the solar-induced chlorophyll fluorescence (SIF) retrieval.

Figure 15 shows an example of a spatial illumination variation within a single ABO2 frame. Footprints 4 and 5 show significant differences in both radiance and slope between their left and right boundaries. These large radiance gradients can produce large clocking errors if not corrected. To correct these errors, the clocking correction algorithm employs the following steps:

1. Single-pixel radiance measurements from color slices are used to characterize spatial gradients across each footprint at several wavelengths across the spectral range sampled by the FPA. This is done by calculating ratios of the footprint-averaged radiances computed by shifting the footprint boundaries in one-pixel (row) increments (after correcting bad pixels in each color slice).
2. Radiance ratios from groups of color slices in nearby columns are combined to increase the signal to noise ratio of each radiance ratio.
3. The radiance ratios for each color slice group are then interpolated across the entire spectral range using a piecewise linear function.
4. This piecewise linear function is then used to define a correction function,  $C(f, \lambda)$ , that is multiplied by the radiance returned in each spectral sample (e.g., for each footprint,  $f$ , and each column,  $\lambda$ ).

Because the clocking introduces radiance variations that propagate almost linearly across the wavelength domain of each FPA, the correction function has a sawtooth shape, with discontinuities at each clocking step. This function is normalized to unity across the spectral domain, so that it does not introduce a shift in the mean radiometric calibration for the spectral channel. A detailed, step-by-step description of the clocking correction algorithm is given in the OCO-2 Level 1B algorithm theoretical basis document (Eldering et al., 2015), along with an analysis of its performance for a typical orbit.

Clocking errors are most easily seen in differences between the observed spectra and best-fit synthetic spectra produced as part of the L2 data product generation. Figure 16



**Figure 16.** The impact of clocking errors and the results of the clocking correction algorithm are most easily seen in the spectrally dependent differences between the observed and simulated spectra. **(a)** Residuals for a spatially inhomogeneous scene are shown in black for the ABO2 (top), WCO2 (middle), and SCO2 (bottom) spectrometers. The corrected spectra are shown in red. Radiance discontinuities occur at the locations of single-pixel shifts in the start and stop row indices of the spectral sample. The amplitudes of the discontinuities are substantially reduced by the clocking correction algorithm. **(b)** Residuals for the ABO2 (top), WCO2 (middle), and SCO2 (bottom) are shown for a scene where the residuals are well corrected in the first two channels but inadequately corrected in the SCO2 channel (red circles). The cause for the poor fit in the SCO2 channel is unknown but may be related to the lack of true continuum in the SCO2 channel.

shows spectral residuals in all three spectral channels before (black) and after (red) the clocking correction is applied to a footprint with a relatively large radiance variation. The clocking errors appear as sharp discontinuities prior to the correction but are virtually invisible after the correction. In Fig. 16a, the maximum amplitudes of the corrections are  $\sim 1.8$ ,  $2.9$ , and  $2.1$  % in the ABO2, WCO2, and SCO2 channels, respectively. The goodness of fit, expressed in terms of the ratio of the root-mean-square (rms) residuals to the random error,  $\chi^2$ , was also improved after the correction, but their relatively large post-correction values indicate that the spatial radiance variations may have been introduced by a sub-footprint-scale cloud. For cloud-free scenes over land, clocking corrections are typically smaller than 1 %. Figure 16b shows a case where the clocking correction worked well for the ABO2 and SCO2 channels but under-corrected the SCO2 channel. These scenes are usually screened out by the OCO-2 L2 retrieval algorithm.

#### 6.4 Polarization orientation anomaly

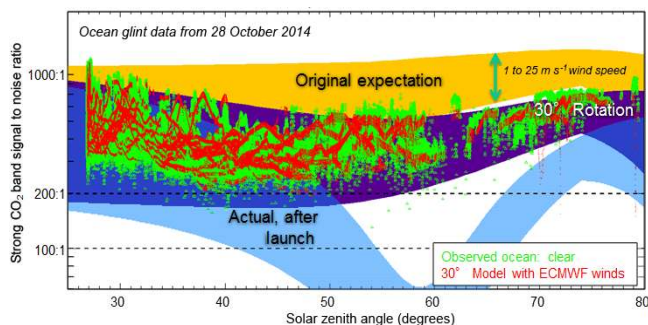
While sunlight incident at the top of Earth's atmosphere is not polarized, scattering by the atmosphere and surface can introduce varying amounts of polarization in the reflected sunlight observed by OCO-2. The degree of polarization can vary substantially but is typically greatest at the Brewster angle over the ocean ( $\sim 53^\circ$  solar incidence angle) or near the terminators, where molecular (Rayleigh) scattering by the atmosphere introduces significant wavelength-dependent polarization. While the degree of polarization is typically unknown, the angle of polarization can be derived from the illumination and observing geometries. For example, if we define a "principal plane" that includes the sun, the surface

footprint, and the instrument aperture, the light polarized parallel to this plane (e.g., the P polarization) is much more strongly attenuated than the light polarized perpendicular to this plane (e.g., the S polarization; Crisp et al., 2008).

Diffraction gratings for high-resolution grating spectrometers, like those used by OCO-2, can also introduce significant polarization in the light that they disperse. 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.

Unfortunately, the polarization angle most efficiently diffracted by the gratings was misinterpreted by the instrument designers. They reported that the instrument was most sensitive to light polarized parallel to the long axis of the slits (Pollock et al., 2010). If this were the case, the largest signals would be obtained if the slits were aligned perpendicular to the principal plane (i.e., parallel to the surface). The spacecraft design and mission observing strategy were optimized to exploit this result (Crisp et al., 2008). In fact, like most grating spectrometers, the OCO and OCO-2 spectrometers are most sensitive to light polarized perpendicular to the long axis of the slits (e.g., in the direction of dispersion).

This design flaw was not caught in the pre-launch instrument assembly and testing program for OCO because the instrument team was focusing on maximizing the throughput of the instrument and was not considering the satellite viewing geometry or polarization characteristics of the surface and atmosphere. This error was not caught for OCO-2 because the instrument and spacecraft were "build-to-print" copies of the original OCO, and similar pre-launch integration and test procedures were followed (Rosenberg et al., 2016).



**Figure 17.** The continuum SNR for ocean glint observations is shown as a function of solar zenith for the SCO2 channel. The gold band shows the performance originally predicted using a Cox–Munk ocean reflectance model for wind speeds varying from 1 to  $25 \text{ m s}^{-1}$ . The actual performance observed immediately after launch, when the spectrometer slits were oriented perpendicular to the principal plane, fell within the light blue band. The low values between 50 and  $60^\circ$  are near the Brewster angle. Predictions from the Cox–Munk ocean reflectance model indicated that introducing a  $30^\circ$  yaw around the center of the instrument’s field of view would yield values within the purple band (actual predictions shown as red points). After this change in the observing geometry was implemented, the observed SNR for clear-sky ocean soundings increased to values above 200 : 1 at most solar zenith angles (green triangles). This is adequate to yield XCO<sub>2</sub> estimates with single-sounding random errors near 1.0 ppm. The SNR in the other two channels is somewhat higher because the solar flux is larger there.

This design error was discovered by the science team almost immediately after launch. It was first seen in the earliest target observations of Total Carbon Column Observing Network (TCCON) sites, where it introduced large viewing-angle-dependent biases in the retrieved properties. These biases vanished if Rayleigh scattering was ignored or if the assumed polarization angle was artificially rotated by  $90^\circ$ . Soon afterwards, the polarization orientation error was seen in the earliest glint observations over the ocean. These observations showed that the radiance levels were dramatically lower than predictions obtained from the ocean surface reflectance model (Cox and Munk, 1956) at solar zenith angles between 50 and  $60^\circ$ , near the Brewster angle for seawater. Individual spectral samples were expected to yield continuum SNR values between 400 and 1000 at these angles, but values much less than 100 were seen (Fig. 17). Note that SNR values near 200 are needed to yield XCO<sub>2</sub> estimates with single-sounding random errors less than 1 ppm.

Once the root cause was identified, a range of mitigation options were studied, and one of the simplest was identified and implemented. Because the OCO-2 spectrometers are the only payload on the spacecraft, the observing geometry was modified to rotate the entire spacecraft by  $30^\circ$  (clockwise from above) about the center of the instrument boresight for glint observations. This corresponds to the “yaw angle” of the spacecraft. With this  $30^\circ$  yaw angle, the SNR values over

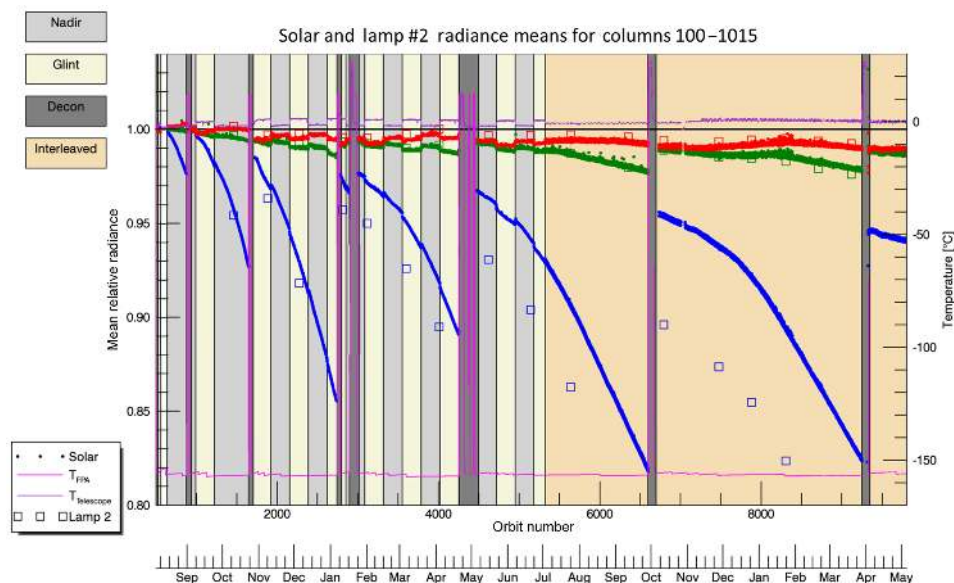
the ocean were restored to values between 200 and 600 at latitudes spanning the Brewster angle (Fig. 17). This yaw angle reduces the amount of sunlight incident on the solar panels, but the power system had ample margin to support this change. The Level 1B products now reflect these modified yaw angle, polarization angle, and footprint-dependent Stokes coefficients. For nadir observations, the spacecraft initially continued to operate with the slit oriented perpendicular to the principal plane because these observations returned adequate signal only over land, where the reflected sunlight is not strongly polarized. On 12 November 2015, a  $30^\circ$  yaw was also adopted for nadir observations to minimize the thermal perturbations introduced to the OBA in switching between nadir and glint observations.

### 6.5 ABO2 sensitivity variations

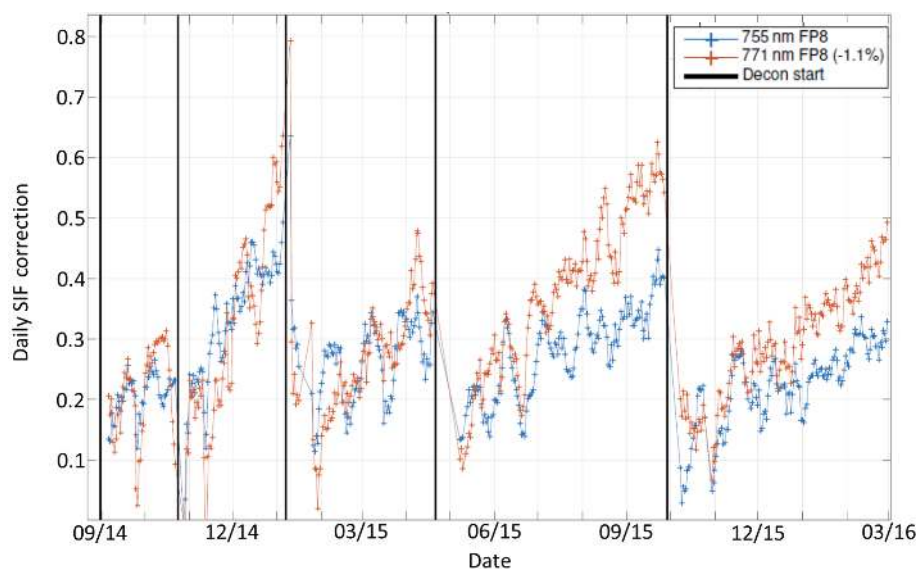
As the radiometric calibration record extended through the first few months of the mission, a trend in the ABO2 channel sensitivity was detected in the routine solar and lamp calibration observations. Little or no change was seen in the CO<sub>2</sub> channels (Fig. 18). Comparisons of solar diffuser observations with the lamp and lunar calibration observations appeared to confirm the basic features of these sensitivity variations. The ABO2 sensitivity appeared to decrease at an accelerating rate while the instrument’s OBA and FPAs were at their operating temperatures. The sensitivity of the ABO2 channel was largely restored after each decon cycle.

An investigation was initiated to identify the root cause of the observed ABO2 signal variations. Both optical contamination and electrical sources were investigated. Water ice accumulation on the FPAs was one of the first optical contaminants considered, but this hypothesis was initially discounted because water ice absorbs much more strongly in the spectral ranges of the CO<sub>2</sub> channels, which showed much less degradation. A very thick layer of ice (many tens of microns) would be needed to produce the observed absorption on the ABO2 FPA. There was no plausible mechanism that would deposit much more water ice (or any other known optical contaminant) in the ABO2 channel than in the CO<sub>2</sub> channels. In addition, a thick layer of ice would significantly alter the ABO2 ILS function, and no changes in the ILS were apparent. A variety of electrical mechanisms for reducing the sensitivity of the FPAs or their signal chains were also considered, but then discarded as tests and analyses progressed. Meanwhile the OCO-2 calibration team developed an empirical correction for the ABO2 sensitivity variations and implemented this correction in the V7 L1B products delivered to the GES DISC and used to generate the V7 L2 products.

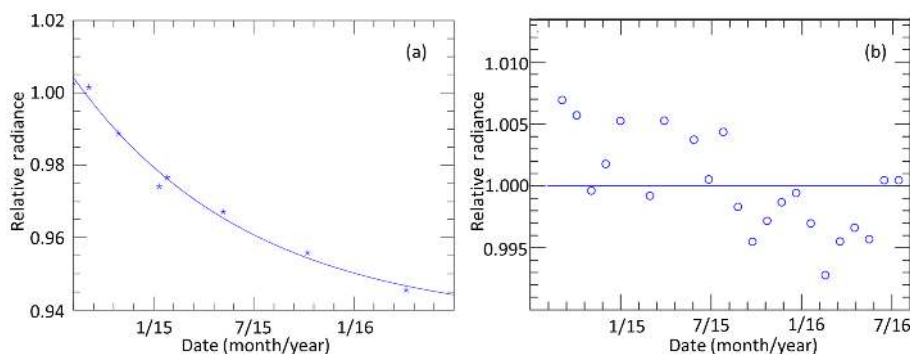
As additional data were collected, other features of the ABO2 degradation were identified. In particular, while the signal degradation was largely restored to its pre-launch value by each decon cycle, the signal did not recover to the pre-launch value. Instead, the recovery level fell slowly over time. This suggested that there may be at least two possi-



**Figure 18.** The relative radiometric response of the ABO2 (blue), WCO2 (green), and SCO2 (red) channels is shown as a function of time (orbit number and month) for the OCO-2 instrument. All data are referenced to the first on-orbit radiometric calibration data, collected in early August 2014. These results show the average response of all footprints in each channel. Routine solar diffuser observations are shown as points, while monthly observations collected with a backup lamp (Lamp 2) are shown as boxes. The instrument decontamination (decon) activities are indicated as dark grey boxes. The alternating 16-day glint and nadir observation periods performed during the first year of the mission are indicated as light yellow and light grey boxes. These observing modes modify the instrument and spacecraft thermal environment, and produced small but measurable changes in the radiometric performance of the instrument and its calibration system. The FPA temperatures are indicated on the right-hand axis. Early in the mission, decon activities were performed frequently to remove ice accumulation on the thermal straps connecting the FPAs to the cryocooler. These activities have become less frequent over time as the system has vented water into space.



**Figure 19.** The correction applied to the solar-induced chlorophyll fluorescence (SIF) estimate is shown as a function of time for the solar lines located near 755 (blue) and 771 nm (red). The timing of the decon activities is indicated by the vertical black lines. These observations suggest that as ice accumulates on the ABO2 FPA, increasing its reflectance, a fraction of that light may be scattered back to the FPA by the cold filter, and other optics above the FPA, producing a zero-level offset.



**Figure 20.** (a) The “slow degradation” in the radiometric response immediately after decon activities (stars), as seen for observations of the Sun through the solar diffuser, can be fit by a smoothly varying function (solid line). The total degradation over the first 18 months in orbit is just over 5 %. Approximately 80 % of this degradation has been attributed to degradation of the solar diffuser plate itself, which is not in place during science observations. (b) Observations of the ~ 75 % gibbous moon taken with the instrument aperture open (no diffuser) show a decrease in the relative response of the ABO2 channel that is less than 1 % yr<sup>-1</sup>. These lunar calibration observations indicate that the remaining ~ 4 % of slow degradation can be attributed to reductions in the gold coatings used on the solar diffuser.

bly independent sources for the ABO2 sensitivity variations. The first resulted in a rapid signal decay that accelerated with time but was largely corrected when the instrument OBA and FPAs were warmed to near room temperature during decon activities. The second was slow and monotonic, and was not affected by the decon cycles. These “fast” and “slow” contributions to the sensitivity degradation were studied in parallel.

### 6.5.1 The fast degradation

As the investigation proceeded, the discovery of what appeared to be ice crystals forming on the ABO2 FPA during pre-launch testing once again brought further scrutiny to this mechanism as a possible source of the fast degradation. A plausible mechanism was soon identified. The ice was not reducing the sensitivity of the ABO2 FPA by absorbing the incident light but by degrading the anti-reflection (AR) coating on the FPA. The ABO2 FPA is made of silicon, which has a relatively high refractive index ( $\sim 3.7$ ) and a high first-surface reflectance ( $\sim 30\%$ ) at the wavelengths covered by this channel. To improve its sensitivity, the vendor had added a thin ( $\sim 120$  nm thick) AR coating to the front surface of the FPA, reducing its reflectance to  $< 10\%$ . The AR coating has an index of refraction that is similar to that of water ice. As ice accumulated on the FPA, and its thickness approached that of the AR coating, the surface reflectance increased at an increasing rate. When the FPAs were warmed to room temperature during a decon cycle, the ice sublimated, and the performance of the AR coating was restored. This mechanism explains the observed depth and rate of the fast component of the sensitivity degradation, and its recovery.

The very thin ice coating on the ABO2 FPA (less than 1/180th the width of a pixel) helps to explain why the fast degradation was not accompanied by a significant change in the ABO2 ILS or other aspects of the instrument performance. More recently, a small zero-level offset (i.e., an offset

with respect to the calibrated dark offset level), with an amplitude as large as 0.8 % of the continuum, has been seen when the ABO2 signal degradation is at its maximum. This change was first detected in the core-to-continuum ratios of the solar Fraunhofer lines used to measure SIF (Fig. 19). It has since been seen in the routine lamp and lunar calibration observations. This zero-level offset probably occurs as some of the light reflected by the degraded AR coating on the ABO2 FPA is reflected back to the FPA by the spectrometer optics. It is currently being characterized and will be corrected in future versions of the OCO-2 L1B product.

This mechanism for the fast degradation provides some insight into why much smaller changes are seen in the CO<sub>2</sub> channels. The HgCdTe FPAs used in the WCO2 and SCO2 channels have a lower intrinsic surface reflectance and use thicker, higher index AR coatings than those used in the ABO2 channel. A much thicker layer of ice is therefore needed to degrade the performance of the AR coatings of these channels. A small amount of degradation has been detected in the sensitivity of the WCO2 channel (see the green line in Fig. 18) that may also be caused by this mechanism.

### 6.5.2 The slow degradation

While a thin layer of ice accumulating on the ABO2 FPA provides a plausible mechanism for the fast degradation, it does not explain the slow, monotonic ABO2 signal loss seen in the solar diffuser and lamp measurements, which is not corrected by the instrument decon activities. The solar calibration observations indicate that the amplitude of this contribution exceeded 5 % during the first 18 months of the mission, but the rate of change seems to be decreasing with time (Fig. 20a). An in-depth analysis of lunar calibration observations, however, indicates a much smaller rate of signal loss than that inferred from the solar and lamp data (Fig. 20b). These measurements, which are made in the normal Earth-

observing mode, without the diffuser in place, indicate that only about one-fifth of the observed attenuation can be attributed to reductions in the throughput of the telescope and spectrometers. Most ( $\sim 80\%$ ) of the apparent slow degradation appears to be caused by a reduction in the throughput of the solar diffuser, perhaps associated with degradation of the gold coatings on its active inner surfaces. The remaining  $\sim 1\%$  decrease in signal over the first 18 months in orbit is most likely due to degradation of the gold coatings in the telescope mirrors and other optical components. This source of signal loss was anticipated prior to launch and was expected to have its largest impact on the ABO2 channel. However, its actual rate was impossible to predict.

The gold coatings on the lamp diffuser (on the back of the aperture cover) are much more exposed to the space environment than those in the solar diffuser and spectrometer optics, and they show larger reductions in reflectivity since launch (see square boxes in Fig. 18). Comparisons of results obtained using the primary calibration lamp, which is used on all nominal polar calibration orbits (Lamp 1), to that of the other two lamps, which are used less frequently (Lamps 2 and 3), indicate that the output of this calibration lamp has also decreased somewhat in the ABO2 channel. Fortunately, the rate of decay of the reflectance of these coatings and the lamps appears to be decreasing with time and can be adequately characterized with lunar observations and observations of well-characterized vicarious calibration sites.

While these processes are correctable and do not threaten mission success, they do reduce the accuracy of the absolute radiometric calibration of the V7 L1B products. The initial ABO2 degradation correction that was applied to this product did not discriminate between the fast and slow degradation. This empirical correction assumed that both the fast and slow components of the signal loss were due to changes in the throughput of the instrument's science optical path or the sensitivity of its FPAs. By including the slow degradation of the solar diffuser in this correction, the radiometric calibration process has over-corrected L1B products since launch. This error has recently been confirmed by lunar calibration data and in observations of parts of the Sahara and Arabian Desert, which appear to be brightening over time in the OCO-2 product, relative to nearly coincident results collected by the Moderate Resolution Imaging Spectroradiometer on the NASA Aqua platform (MODIS/Aqua). This error is currently being more fully characterized using lunar observations and vicarious calibration observations of Railroad Valley, NV, and through additional comparisons with observations collected by MODIS/Aqua and the Japanese Greenhouse gases Observing SATellite (GOSAT, nicknamed "Ibuki"). It will be corrected in future versions of the OCO-2 L1B product.

## 7 Conclusions

OCO-2 was successfully launched on 2 July 2014. Two months later, it began routinely returning almost 1 million soundings over the sunlit hemisphere each day. While its three-channel imaging grating spectrometer was extensively characterized and calibrated prior to launch (Frankenberg et al., 2015; Rosenberg et al., 2016; Lee et al., 2016) and is returning high-quality data, a number of subtle calibration challenges were identified during its first year and a half in orbit. To address these issues, the instrument calibration and observing strategy have been continuously refined during this period.

Starting in June 2015, the OCO-2 team began reprocessing the entire OCO-2 data record, extending back to 6 September 2014, using the V7/7r algorithm and delivering this product to the GES DISC for distribution to the science community. This product includes calibrated, geolocated spectral radiances (L1B products), and retrieved geophysical quantities, including spatially resolved estimates of XCO<sub>2</sub>, surface pressure, and solar-induced chlorophyll fluorescence (L2 products). Updates for bad pixels, cosmic rays, clocking errors, and some aspects of the ABO2 sensitivity variations have been incorporated into the V7 L1B product. The impacts of a polarization anomaly were mitigated through changes in the measurement approach. Other issues, such as thermally induced changes in the instrument dispersion, are not being corrected in the L1B product because a dispersion correction is a routine part of the L2 processing. These corrections will be incorporated into future versions of the L1B product. Corrections for other calibration issues, such as the zero-level offset associated with the ABO2 fast degradation and the slow degradation of the solar diffuser, are under development and will also be incorporated into future L1B products.

These updates to the instrument calibration have fully exploited the onboard calibration system, as well as observations of the Sun; the Moon; and well-calibrated "vicarious calibration" targets on the Earth's surface, such as Railroad Valley, NV. Comparisons with nearly coincident observations collected by MODIS/Aqua and GOSAT are also being used to diagnose and reduce calibration uncertainties. The impact of these uncertainties on the L2 products is being evaluated using comparisons with observations from the TCCON network and other standards. Using these comparisons, a bias correction has been developed and delivered to the community in the V7 "Lite Files" (Eldering et al., 2016). With this correction, OCO-2 usually returns XCO<sub>2</sub> estimates with single-sounding random errors and biases between 0.5 and 1 ppm at solar zenith angles as large as 70°. The median difference between OCO-2 and TCCON XCO<sub>2</sub> estimates is less than 0.5 ppm, and the rms differences are typically < 1.5 ppm (Wunch et al., 2016).

## 8 Data availability

As fully described in Sect. 4 of this paper, the results presented here are based on NASA OCO-2 data products (including relevant metadata and documentation), which are publicly available through the NASA Goddard Earth Science Data and Information Services Center (GES-DISC), at the website <https://disc.gsfc.nasa.gov/OCO-2>.

**Acknowledgements.** Part of the research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The CSU/CIRA contribution to this work was supported by JPL subcontract 1439002. Government sponsorship is acknowledged.

Edited by: A. Kokhanovsky

Reviewed by: four anonymous referees

## References

- Bai, Y., Bajaj, J., Beletic, J. W., Farris, M. C., Joshi, A., Lauxtermann, S., Petersen, A., and Williams, G.: Teledyne Imaging Sensors: Silicon CMOS imaging technologies for x-ray, uv, visible and near infrared, *Proc. SPIE*, 7021, doi:10.1117/12.792316, 2008.
- Beletic, J. W., Blank, R., Gulbransen, D., Lee, D., Loose, M., Piquette, E. C., Sprafke, T., Tennant, W. E., Zandian, M., and Zino, J.: Teledyne Imaging Sensors: Infrared imaging technologies for Astronomy & Civil Space, *Proc. SPIE*, 7201, doi:10.1117/12.790382, 2008.
- Bösch, H., Toon, G., Sen, B., Washenfelder, R., Wennberg, P., Buchwitz, M., de Beek, R., Burrows, J., Crisp, D., Christi, M., Connor, B. J., Natraj, V., and Yung, Y. L.: Space-based near-infrared CO<sub>2</sub> measurements: Testing the Orbiting Carbon Observatory retrieval algorithm and validation concept using SCIAMACHY observations over Park Falls, Wisconsin, *J. Geophys. Res.*, 111, 148–227, doi:10.1029/2006JD007080, 2006.
- Bösch, H., Baker, D., Connor, B., Crisp, D., and Miller, C.: Global Characterization of CO<sub>2</sub> Column Retrievals from Shortwave-Infrared Satellite Observations of the Orbiting Carbon Observatory-2 Mission, *Remote Sens.*, 3, 270–304, 2011.
- Connor, B. J., Bösch, H., Toon, G., Sen, B., Miller, C. E., and Crisp, D.: Orbiting carbon observatory: Inverse method and prospective error analysis, *J. Geophys. Res.*, 113, D05305, doi:10.1029/2006JD008336, 2008.
- Connor, B., Bösch, H., McDuffie, J., Taylor, T., Fu, D., Frankenberg, C., O'Dell, C., Payne, V. H., Gunson, M., Pollock, R., Hobbs, J., Oyafuso, F., and Jiang, Y.: Quantification of uncertainties in OCO-2 measurements of XCO<sub>2</sub>: simulations and linear error analysis, *Atmos. Meas. Tech.*, 9, 5227–5238, doi:10.5194/amt-9-5227-2016, 2016.
- Cox, C. and Munk, W.: Slopes of the sea surface deduced from photographs of sun glitter, *Bull. Scripps Inst. Oceanogr.*, 6, 401–488, 1956.
- Crisp, D., Atlas, R. M., Breon, F.-M., Brown, L. R., Burrows, J. P., Ciaia, P., Connor, B. J., Doney, S. C., Fung, I. Y., Jacob, D. J.,

## D. Crisp et al.: The on-orbit performance of the OCO-2

- Miller, C. E., O'Brien, D., Pawson, S., Randerson, J. T., Rayner, P., Salawitch, R. J., Sander, S. P., Sen, B., Stephens, G. L., Tans, P. P., Toon, G. C., Wennberg, P. O., Wofsy, S. C., Yung, Y. L., Kuang, Z., Chudasama, B., Sprague, G., Weiss, B., Pollock, R., Kenyon, D., and Schroll, S.: The Orbiting Carbon Observatory (OCO) mission, *Adv. Space Res.*, 34, 700–709, 2004.
- Crisp, D., Miller, C. E., and DeCola, P. L.: 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, 2008.
- Crisp, D., Fisher, B. M., O'Dell, C., Frankenberg, C., Basilio, R., Bösch, H., Brown, L. R., Castano, R., Connor, B., Deutscher, N. M., Eldering, A., Griffith, D., Gunson, M., Kuze, A., Mandrake, L., McDuffie, J., Messerschmidt, J., Miller, C. E., Morino, I., Natraj, V., Notholt, J., O'Brien, D. M., Oyafuso, F., Polonsky, I., Robinson, J., Salawitch, R., Sherlock, V., Smyth, M., Suto, H., Taylor, T. E., Thompson, D. R., Wennberg, P. O., Wunch, D., and Yung, Y. L.: The ACOS CO<sub>2</sub> retrieval algorithm – Part II: Global XCO<sub>2</sub> data characterization, *Atmos. Meas. Tech.*, 5, 687–707, doi:10.5194/amt-5-687-2012, 2012.
- Crisp, D. for the OCO-2 Team: Measuring Atmospheric Carbon Dioxide from Space with the Orbiting Carbon Observatory-2 (OCO-2), *Proc. SPIE*, 9607, 960702, doi:10.1117/12.2187291, 2015.
- Eldering, A., Pollock, R., Lee, R. A. M., Rosenberg, R., Oyafuso, F., Crisp, D., Chapsky, L., and Granat, R.: Orbiting Carbon Observatory (OCO) – 2 Level 1B Theoretical Basis Document, available at: [http://disc.sci.gsfc.nasa.gov/OCO-2/documentation/oco-2-v7/OCO2\\_L1B\\_ATBD.V7.pdf](http://disc.sci.gsfc.nasa.gov/OCO-2/documentation/oco-2-v7/OCO2_L1B_ATBD.V7.pdf) (last access: 16 June 2016), 2015.
- Eldering, A., O'Dell, C. W., Wennberg, P. O., Crisp, D., Gunson, M. R., Viatte, C., Avis, C., Braverman, A., Castano, R., Chang, A., Chapsky, L., Cheng, C., Connor, B., Dang, L., Doran, G., Fisher, B., Frankenberg, C., Fu, D., Granat, R., Hobbs, J., Lee, R. A. M., Mandrake, L., McDuffie, J., Miller, C. E., Myers, V., Natraj, V., O'Brien, D., Osterman, G. B., Oyafuso, F., Payne, V. H., Pollock, H. R., Polonsky, I., Roehl, C. M., Rosenberg, R., Schwandner, F., Smyth, M., Tang, V., Taylor, T. E., To, C., Wunch, D., and Yoshimizu, J.: The Orbiting Carbon Observatory-2: First 18 months of Science Data Products, *Atmos. Meas. Tech. Discuss.*, doi:10.5194/amt-2016-247, in review, 2016.
- 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.: The Orbiting Carbon Observatory (OCO-2): spectrometer performance evaluation using pre-launch direct sun measurements, *Atmos. Meas. Tech.*, 8, 301–313, doi:10.5194/amt-8-301-2015, 2015.
- Haring, R., Pollock, R., Sutin, B. M., and Crisp, D.: The Orbiting Carbon Observatory instrument optical design, *Proc. SPIE*, 5523, 51–62, 2004.
- Haring, R. E., Pollock, R., Sutin, B. M., Blakley, R., Scherr, L. M., and Crisp, D.: Fabrication and Assembly Integration of the Orbiting Carbon Observatory Instrument, *Infrared Spaceborne Remote Sensing and Instrumentation XVI*, edited by Marija Strojnik, *Proc. SPIE*, 7082, 708213, doi:10.1117/12.796289, 2008.
- Lamborn, B. B.: Design and Performance of the Cryogenic Subsystem for the Orbiting Carbon Observatory, *Cryogenics*, 48, 198–205, 2008.

- L'Ecuyer, T. S. and Jiang, J. H.: Touring the atmosphere aboard the A-Train, *Phys. Today*, 63, 36–41, doi:10.1063/1.3463626, 2010.
- Lee, R. A. M., O'Dell, C., Wunch, D., Roehl, C. M., Osterman, G. B., Blavier, J.-F., Rosenberg, R., Chapsky, L., Frankenberg, C., Hunyadi-Lay, S. L., Fisher, B. M., Rider, D. M., Crisp, D., and Pollock, R.: Preflight Spectral Calibration of the Orbiting Carbon Observatory, *IEEE T. Geosci. Remote Sens.*, submitted, 2016.
- Miller, C. E., Crisp, D., DeCola, P. L., Olsen, S. C., Randerson, J. T., Michalak, A. M., Alkhaled, A., Rayner, P., Jacob, D. J., Suntharalingam, P., Jones, D. B. A., Denning, A. S., Nicholls, M. E., Doney, S. C., Pawson, S., Bösch, H., Connor, B. J., Fung, I. Y., O'Brien, D., Salawitch, R. J., Sander, S. P., Sen, B., Tans, P., Toon, G. C., Wennberg, P. O., Wofsy, S. C., Yung, Y. L., and Law, R. M. L.: Precision requirements for space-based XCO<sub>2</sub> data, *J. Geophys. Res.* 112, D10314, doi:10.1029/2006JD007659, 2007.
- Na-Nakornpanom, A., Johnson, D., and Naylor, B.: OCO-2 cryocooler development, integration and test, 18th International Cryocooler Conference, Syracuse, New York, 9–12 June 2014, Boulder: ICC Press, 105–114, 2015.
- O'Dell, C. W., Connor, B., Bösch, H., O'Brien, D., Frankenberg, C., Castano, R., Christi, M., Eldering, D., Fisher, B., Gunson, M., McDuffie, J., Miller, C. E., Natraj, V., Oyafuso, F., Polonsky, I., Smyth, M., Taylor, T., Toon, G. C., Wennberg, P. O., and Wunch, D.: The ACOS CO<sub>2</sub> retrieval algorithm – Part 1: Description and validation against synthetic observations, *Atmos. Meas. Tech.*, 5, 99–121, doi:10.5194/amt-5-99-2012, 2012.
- Pollock, H. R., Haring, R. E., Holden, J. R., Johnson, D. L., Kaptanoff, A., Mohlman, D., Phillips, C., Randall, D., Rechsteiner, D., Rivera, J., Rodriguez, J. I., Schwuchert, M. A., and Sutin, B. M.: The Orbiting Carbon Observatory Instrument: Performance of the OCO Instrument and Plans for the OCO-2 Instrument, *Proc. SPIE 7826, Sensors, Systems, and Next-Generation Satellites XIV*, 78260W, doi:10.1117/12.865243, 2010.
- Rauscher, B. J., Stahle, C., Hill, R. J., Greenhouse, M., Beletic, J., Babu, S., Blake, P., Cleveland, K., Cofie, E., Eegholm, B., Engelbracht, C. W., Hall, D. N. B., Hoffman, A., Jeffers, B., Jhabvala, C., Kimble, R. A., Kohn, S., Kopp, R., Lee, D., Leidecker, H., Lindler, D., McMurray Jr., R. E., Misselt, K., Mott, D. B., Ohl, R., Pipher, J. L., Piquette, E., Polis, D., Pontius, J., Rieke, M., Smith, R., Tennant, W. E., Wang, L., Wen, Y., Willmer, C. N. A., and Zandian, M.: JWST near-infrared detector degradation – finding the problem, fixing the problem, and moving forward, *AIP Adv.*, 2, 021901, doi:10.1063/1.4733534, 2012.
- Rayner, P. J. and O'Brien, D. M.: The utility of remotely sensed CO<sub>2</sub> concentration data in surface source inversions, *Geophys. Res. Lett.*, 28, 175–178, 2001.
- Rosenberg, R., Maxwell, S., Johnson, B. C., Chapsky, L., Lee, R. A. M., and Pollock, R.: Preflight Radiometric Calibration of Orbiting Carbon Observatory 2, *IEEE T. Geosci. Remote Sens.*, in press, 2016.
- Stassinopoulos, E. G. and Raymond, J. P.: The Space Radiation Environment for Electronics, *Proc. IEEE*, 76, 1423–1442, 1988.
- Taylor, T. E., O'Dell, C. W., Frankenberg, C., Partain, P., Cronk, H. Q., Savtchenko, A., Nelson, R. R., Rosenthal, E. J., Chang, A. Y., Fisher, B., Osterman, G., Pollock, R. H., Crisp, D., Eldering, A., and Gunson, M. R.: Orbiting Carbon Observatory-2 (OCO-2) cloud screening algorithms; validation against collocated MODIS and CALIOP data, *Atmos. Meas. Tech.*, 9, 973–989, doi:10.5194/amt-9-973-2016, 2016.
- Wunch, D., Wennberg, P. O., Osterman, G., Fisher, B., Naylor, B., Roehl, C. M., O'Dell, C., Mandrake, L., Viatte, C., Griffith, D. W., Deutscher, N. M., Velasco, V. A., Notholt, J., Warneke, T., Petri, C., De Maziere, M., Sha, M. K., Sussmann, R., Rettinger, M., Pollard, D., Robinson, J., Morino, I., Uchino, O., Hase, F., Blumenstock, T., Kiel, M., Feist, D. G., Arnold, S. G., Strong, K., Mendonca, J., Kivi, R., Heikkinen, P., Iraci, L., Podolske, J., Hillyard, P. W., Kawakami, S., Dubey, M. K., Parker, H. A., Sepulveda, E., Rodriguez, O. E. G., Te, Y., Jeseck, P., Gunson, M. R., Crisp, D., and Eldering, A.: Comparisons of the Orbiting Carbon Observatory-2 (OCO-2) XCO<sub>2</sub> measurements with TCCON, *Atmos. Meas. Tech. Discuss.*, doi:10.5194/amt-2016-227, in review, 2016.