

Supplementary Material for The Airborne Chicago Water Isotope Spectrometer: An Integrated Cavity Output Spectrometer for Measurements of the HDO/H₂O Isotopic Ratio in the Asian Summer Monsoon

Benjamin W. Clouser^{1,2}, Laszlo C. Sarkozy², Clare E. Singer^{2,3}, Carly C. KleinStern¹, Adrien Desmoulin², Dylan Gaeta², Sergey Khaykin⁴, Stephen Gabbard⁵, Stephen Shertz⁵, and Elisabeth J. Moyer²

¹Department of Physics, University of Chicago, Chicago, IL, USA

²Department of the Geophysical Sciences, University of Chicago, Chicago, IL, USA

³now: Lamont-Doherty Earth Observatory, Columbia University, Palisades, NY, USA

⁴LATMOS/IPSL, UVSQ, Sorbonne Université, CNRS, Guyancourt, France

⁵National Center for Atmospheric Research, Boulder, CO, USA

Correspondence: Benjamin W. Clouser (bclouser@uchicago.edu); Elisabeth Moyer (moyer@uchicago.edu)

1 Region Choice for Flow Rate Estimation

To estimate the flow rate through the ChiWIS instrument, we chose a flight leg with high natural variability. The selected interval is from the fourth flight of the ACCLIP 2021 campaign on August 17th. This interval features several large changes in mixing ratio approaching 5 ppmv in 5 seconds. Figure S1 shows 62 Hz ChiWIS data (black) and 20 Hz DLH data (cyan).

2 Optical Pedestal

2.1 Pedestal Characteristics

The pedestal radiation originates on or near the laser diode, as evidenced by the fact that it is well-collimated. Masking the first spot on the output mirror reduces the pedestal power by about 60%, implying that the typical mirror reflectivity is about 65% at the characteristic pedestal wavelength. Based on the mirror reflectivity curve, this radiation has a wavelength farther from the laser line center than about ± 300 nm.

Tests with filters centered at other nearby wavelengths showed varying degrees of attenuation, suggesting a broadband, incoherent, non-thermal source, centered around the laser wavelength. A filter centered at 2750 nm (Thorlabs p/n FB2750-500) blocked nearly all of the broadband radiation.

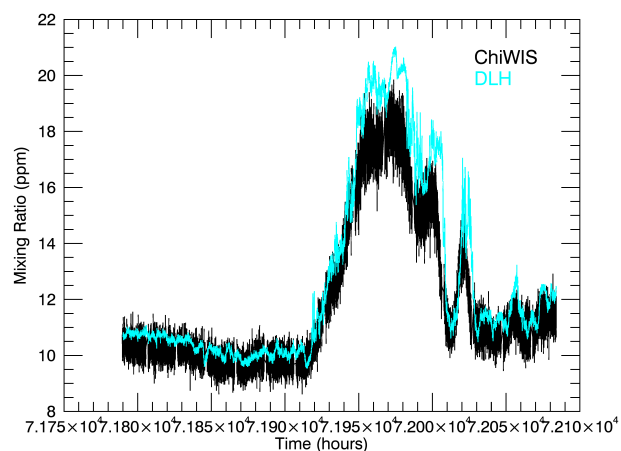


Figure S1. The subinterval of data used to generate the estimation of cavity flush time. Data are taken from the ACCLIP flight on August 17th, 2021. Rapid changes of ≈ 5 ppmv occur within about 5 seconds, allowing for a test of the cavity flush time of the ChiWIS instrument. ChiWIS data are reported at 62 Hz and DLH data are reported at 20 Hz.

To establish the pedestal's linearity, we slowly flushed the cavity with dry air. Line centers that do not vary over a large range of N_2O values are considered optically saturated. From this set of points, we generate a linear function to describe the observed pedestal power. By slightly changing the temperature set point of the laser, we can shift the saturated line centers up and down the ramp, and verify that the pedestal is very nearly linear.

To test the effect of N_2O absorption on the pedestal, we conducted experiments in which the cell volume was rapidly switched between pure N_2O and dry air. During the dry air intervals, the FB2750-500 filter was rapidly switched in and out of the beamline. Using adjacent unfiltered and filtered scans, we reconstructed the pedestal by modeling the unfiltered data as the sum of a linear function and the filtered data multiplied by a constant, where the linear component represents the pedestal and the constant accounts for the attenuation of the main beam by the filter. See the supplementary material for more details on this procedure. Pedestals derived using this method were in excellent agreement with those derived from the N_2O method.

2.2 Pedestal Correction Procedure

Another phenomenon made correcting the spectral data for the pedestal effect even harder: the temperature dependence of the detector sensitivity. Therefore, variations in detector temperature caused variations in the pedestal height. We derive a strong, nearly linear, correlation between observed power and the detector mount temperature to account for in-flight variations in detector temperature and sensitivity. In the StratoClim configuration, the detector was not sufficiently anchored to a heat sink, meaning that over the course of a flight it would warm up by several degrees, resulting in decreased sensitivity. The photodiode temperature was not recorded during StratoClim, but the detector mount temperature, which was recorded, is an effective proxy.

The laboratory data used to derive this relationship are shown in Figure S2 (SigData_RV). The black dots represent measurements of the pedestal height at the center of the big water line (3777.9492 cm^{-1} , roughly sample number 1200 in a typical scan) as the detector head warms from about 26.5 to 40.5 degrees C, spanning the range of detector mount temperatures observed during StratoClim. In the following, we derive a linear scaling between detector mount temperature and the pedestal slope and intercept.

SigData_FO provides a canonical slope and y-intercept for the pedestal during flight, since that data set was taken in September 2017, immediately after the StratoClim campaign. However, SigData_FO does not span the range of detector mount temperatures observed during the campaign (only 28-30° C), and therefore cannot be used to generate a pedestal correction on its own. SigData_RV spans the range of detector mount temperatures, and linear fits to the pedestal slope vs. detector mount temperature and pedestal y-intercept vs. de-

tor mount temperature relationships form the basis of our quantitative pedestal correction. However, SigData_RV was taken in May 2019, and we must admit the possibility that the laser aged during this period and that the pedestal no longer has the same y-intercept or slope. To account for this possibility, we rescale SigData_RV to match SigData_FO over their common temperature range. This rescaling does not affect the temperature dependence, since that is a property of the detector sensitivity under temperature variation.




Figure S2. Comparison of the scaling between pedestal power and detector mount temperature in flight data and in two laboratory data sets (green, black; see text). For this example, pedestal power is determined at the center of the big water line (3777.9492 cm^{-1}). The red line is a fit to this relationship, which is then used to correct flight data. The six colored points show flight data in saturated intervals during flights 2, 3, 4, 6, 7, and 8 during StratoClim, and the blue line is a fit to those points. The blue and red lines show near identical slopes, validating the use of laboratory data to determine the relationship. Their intercepts differ in magnitude by $\sim 2\%$.

Comparison with In-flight Water

During brief intervals at the start of each StratoClim flight, the instrument samples very wet inlet air which optically saturates the big water line, which allows for the estimation of a pedestal upper bound. The colored dots plotted on Figure S2 correspond to the detector mount temperature and observed power at the tip of the big water line during that flight's high water interval. The blue line is a line of best fit through these points. The slope of that line is not statistically different from that of the red line fit through the lab data points, although the magnitude is greater by about 2%.

This estimate was expected to be higher than, rather than equal to the pedestal due to: a) the difficulty of achieving 'true' optical saturation in ICOS instruments, and b) the transient nature of the very wet time intervals. ICOS output has contributions from radiation that has made very few passes through the optical cavity and therefore experienced little absorption. Simulated ICOS spectra estimate that the tip of the water line would achieve minimum transmission of $\approx 1.1\%$ at 1000 ppmv water. Also, the very wet intervals last only several

seconds, comparable to the averaging time, making it possible that unsaturated scans are averaged into the relevant spectra.

2.3 Test with In-flight Tracers

In addition to H₂O and HDO, ChiWIS also measures several very weak CO₂, N₂O, and CH₄ transitions. These can be used to test the quality of the pedestal correction procedure, using both a priori knowledge about the typical concentrations of a tracer, and through direct comparison to measurements made by other instruments (often with much higher precision) aboard the Geophysica.

We test the quality of the pedestal correction for fitting two very weak N₂O transitions present in the typical ChiWIS spectral regions. These fits on the corrected data pass several consistency checks. First, the retrieved concentrations are consistent with each other to within the HITRAN linestrength error of ± 2 -5%. Given that the N₂O spectral features are spaced quite far apart on the ramp, the presence of the pedestal can introduce a significant disparity between concentrations retrieved from different transitions of the same molecule.

Table S1. Spectroscopic Parameters of ChiWIS N₂O Lines

ν (cm ⁻¹)	S	ΔS	Δn_{air}
3777.7730	1.085e-22	2-5%	5-10%
3778.3655	9.475e-23	2-5%	5-10%

Second, both features return concentrations consistent with the HAGAR instrument, which measures N₂O. Last, fractional deviation between the ChiWIS N₂O concentrations and HAGAR N₂O concentrations are constant from flight to flight, meaning that the pedestal correction procedure does not introduce flight-to-flight variation.

3 Effects of Rebinning

Raw spectral data contain 50000 samples per scan, and are typically rebinned prior to analysis. This saves space and speeds up fitting of the spectral features. However, rebinning can introduce systematic bias through line shape distortion and by confounding retrievals of ringdown times. Both effects are explored here.

3.1 Lineshape distortion

As the rebin size becomes comparable to the width of the spectral features, rebinning warps the spectral features and results in biased retrievals of concentration. To quantify this effect, we fit a section of data from the ACCLIP flight on 20210810 across a range of values. The results are summarized in Figure S3.

The top panel shows concentrations derived from data rebinned by ten samples. The bottom shows deviations from that data under greater rebinnings. Rebinning-induced errors

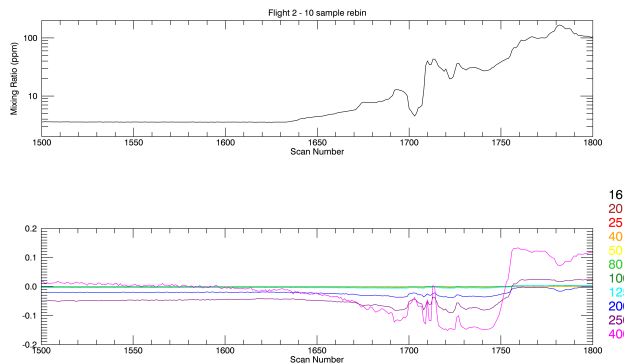


Figure S3. Data from Flight 2 of the 2021 ACCLIP campaign out of Houston on 20210810. The top panel shows H₂O mixing ratio data derived from fits to spectral data rebinned by 10 samples. The bottom panel shows deviations from that mixing ratio under various other rebinnings.

are generally less than about 1% up to 125 samples rebinned, which is roughly comparable to the size of the spectral features in question. Final data are derived from fits to data rebinned by 10 or 20 samples.

3.2 Improper Ringdown Time Retrievals

Accurate ringdown time retrievals are necessary to ensure accurate concentration retrievals from the raw spectra. A typical ringdown scan (black) is shown in Figure S4, along with a fit to the scan (blue). These scans are taken 2-3 times in each flight.

Rebinning mixes samples from before and after the laser turns off, introducing uncertainty into the ringdown time retrieval. This extra uncertainty reduces the constraint on the cutoff time and the ringdown time parameter, resulting in poor ringdown time retrievals. Figure S5 shows the effect of rebinning on retrieved ringdown times. We take the ringdown time retrieved with no rebinning (26.04 microseconds) to be the true value, and the fractional deviation of all subsequent values are referenced to this one. Final data in the StratoClim campaign were generated from data that was rebinned by 10 samples, resulting in an under reporting of ringdown times of less than 0.3%. Ringdown times in the ACCLIP campaign will be generated from raw, unrebbinned data.

4 Data Availability

Failure to return satisfactory data occurred during Flights 1 and 5 of the StratoClim campaign. During Flight 1, a loose contact in the TRB connector carrying the trigger pulse to initiate signal acquisition resulted in a data acquisition rate of about one-tenth of normal. Although the returned spectra are of good quality, to date no effort has been made to analyze

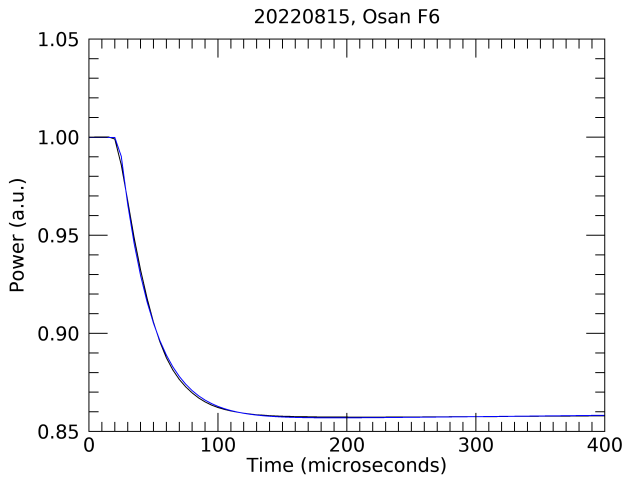


Figure S4. A ringdown scan from the Osan flight on August 15th, 2022. The laser abruptly turns off about 23 microseconds after the trigger signal which initiates data acquisition. The data (black) here are averaged by 5 seconds and rebinned by 20 samples. The fit to the data (blue) yields a ringdown time estimate of 24.9 microseconds.

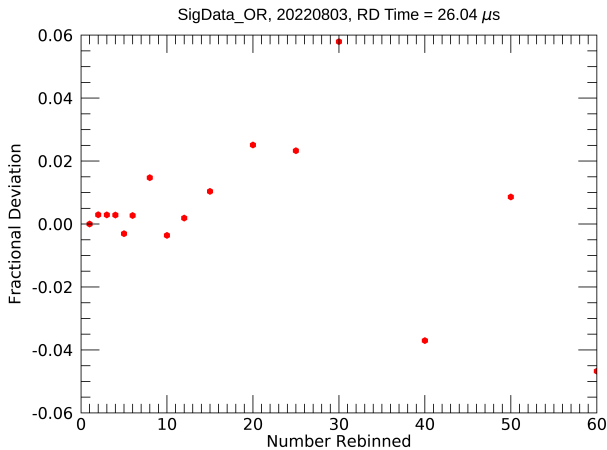


Figure S5. Retrieved ringdown times are sensitive to the number of samples rebinned.

properly analyze this data, partly due to flight through fairly uninteresting air, and partly due to data dropout from other instruments on the flight. This issue was mitigated in the field before Flight 2 by swapping the trigger pulses for ringdown and signal acquisition.

Rolling blackouts in the flight hangar were an issue throughout the StratoClim campaign in Kathmandu. While on ground power before Flight 5 during a pre-flight flush with dry air, one such blackout occurred. Flushing the instrument requires both top and bottom valves to be open, and when the power failed, they both remained open. This allowed wet hangar air to rush into the optical cavity through the bottom valve. Fortunately, the high-reflectivity cavity mirrors were not impacted, and did not require cleaning afterwards. However, data returned from Flight 5 show clear evidence of outgassing throughout the entire flight. To date, no significant effort has been made to analyze or make a correction on data from this flight.

During the third flight of the Summer 2021 ACCLIP campaign, the instrument failed to return satisfactory data due to an issue with the software controlling its top valve. In this case, the new external inlet for the ACCLIP campaign proved to be more difficult to heat than the StratoClim inlet due to the absence of a fairing surrounding it. The under-temperature condition eventually caused the control software to close the top valve. This feature was programmed into the control software to prevent icing within the instrument, but the experience of in-flight conditions during StratoClim has shown that this feature is unnecessary in almost all conditions. Other, internal sources of heat on tubing internal to the instrument are more that sufficient to sublimate any ice particles that are able to enter the inlet before they reach the optical cavity.

During the fourth transit flight west between Adak and Misawa before the ACCLIP 2022 campaign, a computer failure required the restart of the instrument. Some data was lost at the beginning of the flight. During the fifth flight between Misawa and Osan the same issue occurred, and the decision was taken not to power cycle the instrument due to the very short nature of the flight.

At the end of the second flight out of Osan Air Base on 20220804, a DC-DC converter powering the laser driver board failed, resulting in the loss of a small amount of data. The component was replaced and there were no further issues.

No data was recorded on the third Osan flight. The data acquisition computer failed to boot at the beginning of the flight, and did not respond to power cycling. The cause was ultimately determined to be a poor connection between the PXIe computer and its backplane. The board had likely become loose due to vibrations and repeated thermal cycling, and simply reseating the board solved the problem.