JPL Carbon Dioxide Laser Absorption Spectrometer Data Processing Results for the 2010 Flight Campaign

Joseph C. Jacob, Gary D. Spiers, Robert T. Menzies, Lance E. Christensen
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099
USA
{Joseph.Jacob,Gary.Spiers,Robert.Menzies,Lance.Christensen}@jpl.nasa.gov

1. Introduction

As a precursor to and validation of the core technology necessary for NASA's Active Sensing of CO2 Emissions over Nights, Days, and Seasons (ASCENDS) mission, we flew JPL's Carbon Dioxide Laser Absorption Spectrometer (CO2LAS) in a campaign of five flights onboard NASA's DC-8 Airborne Laboratory in July 2010. This is the latest in a series of annual flight campaigns that began in 2006, and our first on the DC-8 aircraft.

Column CO2 concentrations are estimated using the Integrated Path, Differential Absorption (IPDA) technique with near nadir pointing of co-aligned online and offline wavelength channels. At 50 MHz sampling rate, we measure the differential return from ground surface backscatter in the two channels. The frequencies of the online and offline channels are near the center and wings, respectively, of the R30 CO2 absorption line. Heterodyne detection at 2.05-μm wavelength provides high sensitivity with weighting in the lower and middle troposphere.[1]

We have completed data processing of the three flights over the California Central Valley, Mojave Desert, and Pacific Ocean. In each flight we fly several overpasses at stable altitudes from 8 KFt to 40 KFt above the surface. We extract the return signal where it is strongest at the lower altitudes between 8 KFt and 20 KFt. The higher altitudes with minimal return signal, e.g. 40 KFt, are used to compute a baseline noise floor to be removed to aid in return signal power estimation.

In between overpasses, validation measurements are taken by inserting a rotating belt in the laser path to provide a calibration mechanism that is devoid of any significant atmospheric absorption. Comparison of the validation runs taken before and after an overpass with those taken at other times during the flight provides a means to compensate for small but significant variations in transmitter power or overall system gain.

The backscattered return power is estimated via automated detection in the frequency domain. Integration times are carefully managed to balance the need for both adequate signal to noise ratio and platform stability. Longer integration improves signal detection efficiency at the expense of being subject to larger temporal variations in aircraft altitude, attitude, ground elevation, albedo, and atmospheric conditions. Quality control filters can be activated to eliminate data segments from the calculation that are corrupted due to clouds in the field of view, excessive turbulence, or other flight anomalies. We calculate the laser range to ground as a function of time and use it to compensate for variable atmosphere path length during the overpass.

In situ and radiosonde measurements of environmental conditions drive forward model simulations. We compare the differential power return with that predicted by the forward model to compute coumn CO2 concentration as both an overpass average and as a function of time and along track distance.

2. CO2LAS frequency domain processing

Full column CO2 concentration can be estimated by considering an overpass in isolation and comparing with a forward model that extends down to the Earth's surface. Alternatively, the difference in absorption between pairs of overpasses can be used to estimate the CO2 concentration in the column between the boundaries at the two altitudes.

2.1. Periodogram summation strategy

16K-point FFTs are computed to produce periodograms with 3 KHz resolution. We seek to integrate periodograms over long enough time periods to maximize the signal to noise

ratio within reasonable limits of aircraft and environment stability.

Long integration times, e.g., 5 seconds, are used for our backscattered power calculation. Return signal peaks are detected with shorter integration times (e.g., 40 ms). Before we combine the shorter periodogram sums to produce the longer summation, we shift them so that all of the return signal peaks align at their mean. This ensures that the return signal is easily visually identified and that we have adequate bandwidth away from the signal to estimate the baseline level that remains after noise floor flattening, described below.

2.2. Noise floor flattening

We use the 40 KFt altitude overpass with no significant signal return to compute an overpass average noise floor for the offline and online channels. We aim to remove the noise floor from each periodogram because it can mask our true return signal peaks and make accurate peak detection and power estimation more difficult.

The noise floors for both channels are smoothed and then jointly normalized so that the online value at 15 MHz is set to 1. The noise floor component of a periodogram summation for any integration time is then removed by dividing by this smoothed and normalized noise floor for the corresponding channel. This is illustrated in Figure 1 and has the benefits that it flattens the periodogram, and equalizes the peak-to-peak noise level across the spectrum. Both of these make automated signal detection more robust.

2.3. Peak detection

We automatically detect the signal return in both channels at regular intervals during an overpass. Under normal flight conditions we expect to find the return signal in the window from around 8.5 MHz and 19.5 MHz. The actual signal peak center at any time is dependent on the aircraft speed and attitude as indicated by the Doppler equation.

Care is taken during automated return signal detection to avoid mistaking a particularly large noise spike as being the true Doppler return. This is especially important for the case of low signal to noise ratio, which is common at the higher altitudes. We sample the signal in the vicinity of the peak center and require that a majority of the data points meet a threshold multiple of the Root Mean Square (RMS) signal level in a baseline window away from the signal.

2.4. Return power estimation

The return power in each channel is calculated to be the area under the curve in our peak search window for the long periodogram summation after the noise floor flattening is completed and the baseline is subtracted. The return power is scaled to compensate for any deviations of the transmit power from the overpass average.

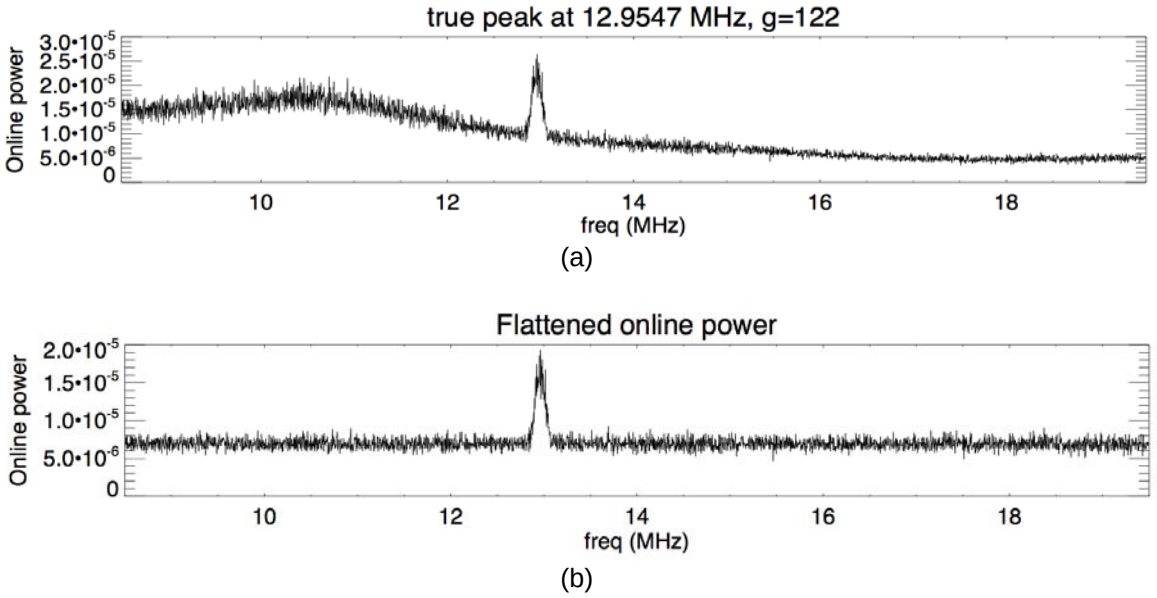

Figure 1. A periodogram with 40 ms integration time. (a) Original periodogram. (b) Periodogram after noise floor flattening. Note that the baseline is flattened and the peak-to-peak noise level is made uniform across the spectrum.

2.4. Range to ground analysis

We compensate for variations in aircraft altitude and attitude, and ground elevation by adjusting return power based on the path length from the laser transmitter to the ground at each instant. Since we do not have a co-boresighted laser altimeter, we compute laser range to ground using aircraft Global Positioning System (GPS) and a Digital Elevation Model (DEM) from the Shuttle Radar Topography Mission (SRTM).

Our laser range to ground algorithm gets the ground elevations from the 1 arc second resolution SRTM version 2.1 ("finished") data product.[2] The NASA DC-8 REVEAL data provides the aircraft altitude, attitude, and speed at 1 second intervals. For flight campaigns where REVEAL data is not available (e.g., for earlier flights on the Twin Otter aircraft over El Mirage and Oklahoma), an Inertial Navigation System (INS) GPS unit is available as part of our instrument package.

To determine which SRTM datum to use for each range to ground calculation, we calculate the intersection point of the laser and the ground. The nadir vector from the laser transmitter to the ground is computed first. A quaternion rotation is applied to account for the spacecraft roll, pitch, and true heading. The intersection point at the ground is then computed by extending this rotated vector to find the intersection with the SRTM topography. We use the return signal Doppler frequency and aircraft speed to calibrate the pitch angle based on the Doppler equation.

Our range to ground analysis revealed that there was significant ground elevation variation across some overpasses. For example, the Mojave/Needles ground track is shown in Figure 2. Therefore, this range to ground analysis is critically needed. Using the SRTM DEM permits us to fly without a laser altimeter, which significantly reduces the cost of our instrument package.

2.5. Quality control filters

Our return power calculations account for normal variations in instrument (e.g., transmitter power), aircraft (e.g., altitude, attitude, velocity), atmosphere (e.g., temperature, pressure), and ground (e.g., elevation, reflectance). However, we have found that our results are more consistent and repeatable if we activate various quality control filters to discard data for short anomalous time periods. These can be indicative of clouds in the field of view, excessive turbulence, or instrument anomalies. We are making progress toward fully automating detection of these phenomena.

We examine the first derivative of the transmitter power time series to find short unstable time segments that are the signature of momentary loss of laser lock. This often causes too much loss in the online channel for detection of the Doppler return. Therefore, we developed software to automatically filter out these time periods from the frequency domain processing data stream.

We can automatically filter out short (i.e. 40 ms) time periods when the Doppler return frequency differs between the offline and online channels by more than a variable threshold amount. In the results reported in this paper, this filter is turned off, but the capability exists for future data processing. We also limit how much the Doppler return frequency can drift in the integration time period for power calculation (i.e. 5 s). As described above, the long integration time for power calculation is first computed as shorter 40 ms summations for peak detection. if any 40 ms time periods has a Doppler return frequency that differs from the mean by greater than a threshold amount (currently set to 2 MHz), it is discarded.

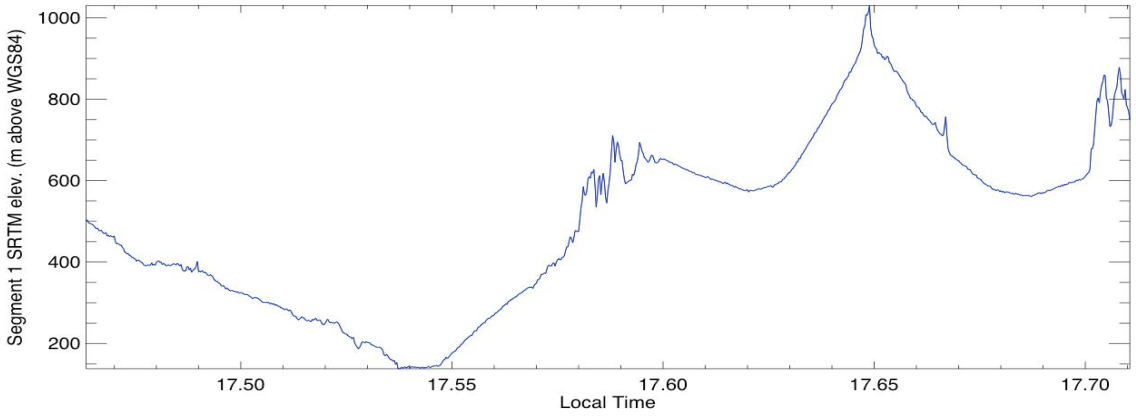

Figure 2. Mojave/Needles ground elevations at laser-ground intersection point as indicated by SRTM.

The long integration times are computed in sequence to make up a time series over a stable overpass. If any of the power values fall more than two standard deviations below the mean, it is discarded.

2.6. Forward modelling

The Differential Absorption Optical Depth (DAOD) is related to both transmittance and return power:

$$DAOD = \ln(\tau_{off} / \tau_{on}) = \tfrac{1}{2} \ln(P_{off} / P_{on}) \quad (1)$$

where τ_{off} and τ_{on} are the offline and online transmittance, and P_{off} and P_{on} are the corresponding return powers. The factor of ½ in Equation 1 is due to the fact that the backscattered signal that reaches the detector has made two passes through the atmosphere: from the transmitter to the Earth, and from the Earth to the detector.

Our current results use the Line By Line Radiative Transfer Model (LBLRTM) from Atmospheric and Environmental Research (AER) for the forward modeling.[2] We define a custom atmosphere with parameters for up to 200 atmosphere layer boundaries. Each layer boundary is assigned a pressure, temperature, and relative humidity from in situ sensors or radiosondes launched from the ground near the middle of the flight track.

For the forward model we assume a fixed column CO_2 density, D, and compute $\ln(\tau_{off} / \tau_{on}) = R_{model}$. We can then compute a time series of measured column CO2 as a function of elapsed flight time:

$$CO_2(t) = D * \ln(P_{off} / P_{on}) / (2 * R_{model}) \quad (2)$$

The factor of 2 in this equation is needed because the forward model log ratio (R_{model}) is for a single pass through the atmospheric column. The aircraft velocity data from REVEAL is used to convert this to column CO2 as a function of distance. An example retrieval result for the California Central Valley 8.5 KFt overpass is shown in Figure 3.

3. Conclusion

Improvements in the data processing algorithms for JPL's CO2LAS instrument are reported with positive results for the Summer 2010 flights. Support for integration times of arbitrary length allow us to carefully balance the need for both high signal to noise ratio and platform stability. The effects of transmit power fluctuations on the return signal power are removed. Range to ground analysis compensates for variations in aircraft altitude and attitude, as well as ground topography. A DEM from SRTM is used for ground elevations, eliminating the need for a co-boresighted laser altimeter. Quality control filters eliminate bias due to short-lived anomalous conditions. Comparison with a forward model allows us to rapidly convert raw flight data into CO2 time series for analysis.

4. Acknowledgements

This research was carried out at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under contract with the National Aeronautics and Space Administration (NASA).

5. References

1. G. D. Spiers, R. T. Menzies, J. Jacob, L. E. Christensen, M. W. Phillips, Y. Choi, and E. V. Browell, Atmospheric CO_2 Measurements with a 2-micrometer Airborne Laser Absorption Spectrometer Employing Coherent Detection, Applied Optics Journal, Vol. 50, Issue 14, pp. 2098-2111, 2011.
2. Farr, T. G., et al. (2007), The Shuttle Radar Topography Mission, Rev. Geophys., 45, RG2004, doi:10.1029/2005RG000183.
3. Clough, S. A., M. W. Shephard, E. J. Mlawer, J. S. Delamere, M. J. Iacono, K. Cady-Pereira, S. Boukabara, and P. D. Brown, Atmospheric radiative transfer modeling: a summary of the AER codes, Short Communication, J. Quant. Spectrosc. Radiat. Transfer, 91, 233-244, 2005.

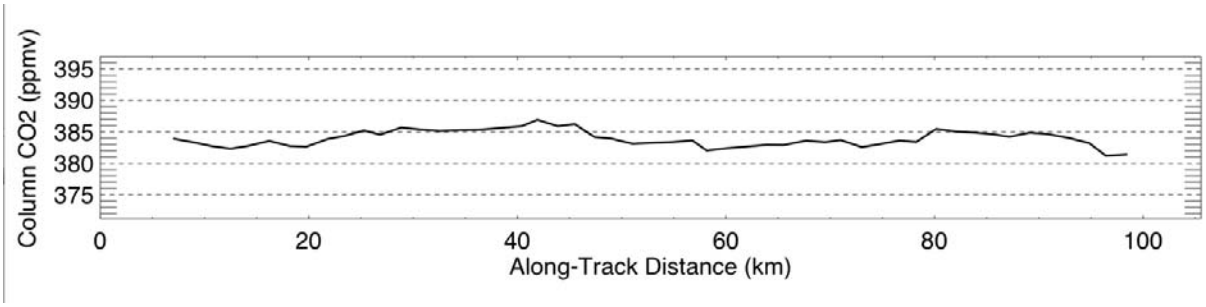

Figure 3. Column CO2 retrieval for the 7/8/2010 8.5 KFt flight over California's Central Valley.