SCISAT-1: Retrieval algorithms, ACE-FTS testing and the ACE database

Ray Nassar, Chris Boone, Kaley A. Walker, Sean D. McLeod and Peter F. Bernath*
Department of Chemistry, University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1

ABSTRACT

The SCISAT-1 mission, also known as the Atmospheric Chemistry Experiment (ACE), is a Canadian satellite mission to investigate the chemical and dynamical processes that control the distribution of ozone in the stratosphere and upper troposphere. The satellite is scheduled to launch in August 2003, carrying two main instruments: a high-resolution infrared Fourier transform spectrometer (ACE-FTS) and a dual grating UV-Vis-NIR spectrometer known as MAESTRO (Measurement of Aerosol Extinction in the Stratosphere and Troposphere Retrieved by Occultation) both operating primarily in solar occultation mode. Aspects of the mission pertaining to work done by ACE science team members from the University of Waterloo will be described, such as: the ACE-FTS forward model for retrieval of temperature, pressure and VMR profiles; ACE-FTS instrument testing and results; and the ACE Database along with data storage and processing hardware.
Keywords: atmospheric remote sensing, Fourier transform spectroscopy, retrieval algorithms, instrument testing.

1. INTRODUCTION

The Atmospheric Chemistry Experiment (ACE) is a Canadian satellite mission to investigate the chemical and dynamical processes that control the distribution of ozone in the stratosphere and upper troposphere. It is scheduled to launch in August 2003 carrying a high-resolution Fourier transform spectrometer (FTS) operating in the 750- 4100 cm^{-1} range and a dual grating spectrometer known as MAESTRO (Measurement of Aerosol Extinction in the Stratosphere and Troposphere Retrieved by Occultation) operating in the 285-1030 nm range. The ACE-FTS is equipped with two solar imagers operating at 1.02 and 0.525 μm. Both the FTS and MAESTRO will operate primarily in solar occultation mode.

ACE will orbit the earth approximately 15 times per day in a circular low-earth orbit of 650 km and an inclination of 74°, enabling it to observe approximately 30 occultation events (sunrises or sunsets) per day. During an occultation, the FTS, MAESTRO and imagers observe the sun through an opening in the baseplate while the FTS passive-cooler points to deep space to cool the FTS detectors. On each occultation, ACE will record a series of atmospheric transmission spectra that will be inverted to give altitude profiles of temperature, pressure and atmospheric constituents.[1] This process involves recording spectra that contain contributions from different altitude regions or layers of the atmosphere. A sunset occultation begins with the FTS recording a solar spectrum with no atmospheric attenuation (an exoatmospheric spectrum), then proceeds by recording atmospheric spectra at a sequence of tangent heights, decreasing by 3-4 km in altitude per measurement, as in Figure 1. Each atmospheric spectrum will be ratioed to the exoatmospheric spectrum to remove solar or instrumental features. The FTS will make measurements of the atmosphere up to 100 km and modeled data will be used to extend retrievals to 150 km, with the retrieved profiles interpolated onto a 1 km grid.

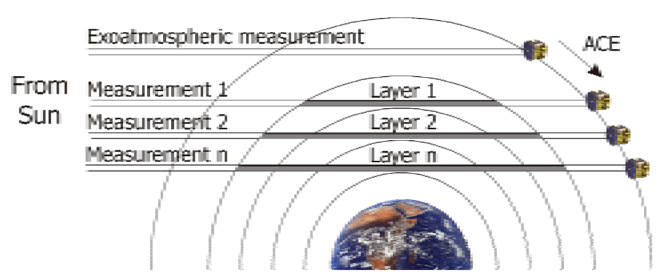

Figure 1: Schematic diagram depicting the satellite motion while recording spectra during a sunset occultation.

*Corresponding author: bernath@uwaterloo.ca; phone 1-519-888-4814.

2. RETRIEVALS

Quantitative interpretation of the atmospheric FTS spectra begins with the calculation of synthetic spectra using what is known as the forward model. In the forward model, a solar ray is traced along the path it would follow through the atmosphere, including the effects of refraction. Using assumed altitude profiles for pressure, temperature and the volume mixing ratios (VMRs) of atmospheric constituents, the transmittance of light as a function of frequency (i.e., the spectrum) over the entire path is then calculated. Forward model calculations are performed on a standard grid, defined by 1 km thick concentric shells, where physical properties (pressure, temperature, VMR) within a shell are assumed constant. Because values of the parameters are only defined at the measurement tangent heights, they must be interpolated onto the standard grid. The VMRs for molecules of interest, pressure and temperature are the parameters in the model that are determined from a nonlinear least squares fit. Only small portions of the spectrum (~0.3 to 1 cm^{-1} wide) containing the most pertinent information are analyzed. These small portions of the spectrum, together with the associated ranges of tangent heights to be included in the analysis, are known as microwindows. The analysis proceeds in two steps, whereby pressure and temperature are determined first, and VMRs for the molecules of interest are determined in a subsequent step.

2.1 Temperature and pressure

For well-resolved FTS spectra, the relative intensities of different lines and the absolute intensities of the lines each provide useful information. Numerous temperature-related factors determine the absorption cross section for a particular line. An important factor is the lower state population, which can vary dramatically from line to line and has an exponential dependence on the energy of the state. Therefore, relative intensities of lines having different lower state energies give a fairly direct measure of temperature. Absolute intensities can be used to determine one additional parameter per measurement, however, multiple parameters (pressure, density, and the VMRs of atmospheric constituents) remain to be determined. Because there is insufficient information to determine all the quantities of interest, constraints must be introduced to reduce the number of independent parameters. For example, by imposing the constraint of the ideal gas law, density can be expressed as a function of pressure and temperature, thereby reducing the number of independent parameters by one. Hydrostatic equilibrium can be introduced as a second constraint, where the most appropriate approach to this depends on the quality of the pointing knowledge provided from the satellite. The quality of pointing knowledge that will be derived from the satellite will not be known for certain until it is in orbit, so algorithms were developed for both the good pointing knowledge and poor pointing knowledge cases. These were described in a previous work [2] but will be summarized here.

2.1.1 Poor pointing knowledge

In the case of poor pointing knowledge, the geometry of the measurement itself becomes another unknown parameter, therefore a molecule for which the VMR is well known must be chosen to be fixed in the analysis as another constraint. The variation in VMR as a function of altitude must be well known for this molecule in order to avoid biasing the results, making CO_2 an excellent candidate for this approach. As shown in Figure 2, CO_2 is perfectly mixed in the middle atmosphere (~15 to ~80 km) resulting in a constant VMR that is well known, and can be fixed for this altitude range.

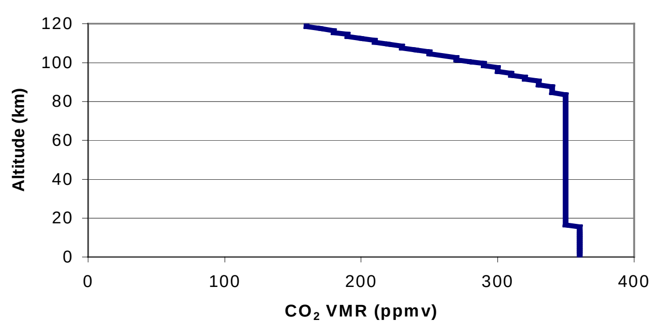

Figure 2: An example of a CO_2 VMR profile from 1994, taken from data archives of the NASA ATMOS missions.[3] The constant VMR portion occurs from about 15-80 km.

Above ~80 km, the variation of CO_2 VMR with altitude is not constant nor well known, so it cannot be fixed. Measurements from the ACE mission will extend up to 100 km in altitude, and for efficiency of processing time, the entire altitude range (10 to 100 km) will be analyzed at once. The altitude range is divided into two regions, high altitudes and low altitudes, where the crossover altitude (the dividing line between the two regions) is taken to be the measurement tangent height nearest to 70 km. When air descends in the winter vortex region, this crossover altitude may need to be chosen somewhat lower. [4]

We assume good pointing knowledge above the crossover point, which is a valid assumption since problems with pointing knowledge (e.g., deteriorated performance of suntracker, refractive distortion of solar image, etc.) are most significant at lower altitudes. For higher tangent heights, good pointing knowledge can be achieved from relatively simple geometry calculations (assuming accurate characterization of the satellite orbit). Thus, for the case of poor pointing knowledge, tangent heights are fixed at high altitudes (to the best values that can be obtained) down to the crossover measurement (the measurement nearest to 70 km). The tangent height for the first measurement below the crossover must also be fixed and the tangent height separations for measurements below the crossover are then calculated, by applying the constraint of hydrostatic equilibrium. The calculation always employs three *successive* measurements, with the tangent height determined in one step required for the next step, working downward from the crossover point to the lowest measurement used in the analysis.

Above the crossover, tangent heights are fixed, pressure is calculated and the parameters to be determined in the least squares fit are 1/T and CO_2 VMR at the tangent heights. For the crossover measurement and below, CO_2 VMR is fixed to its known value, tangent height separations are calculated and the parameters are 1/T and P at the tangent heights. In both cases, two parameters are determined per measurement, the maximum available information content.

2.1.2 Good pointing knowledge

Temperature and pressure retrievals are greatly simplified when good pointing information is available from the satellite. The analysis for high altitudes remains as in the case of poor pointing knowledge, while the analysis procedure for low altitudes differs significantly. At low altitudes, parameters for the least squares fit will again include 1/T at each measurement tangent height. The pressure at one altitude must be included as a parameter in the fit or explicitly fixed. The ACE analysis software takes the pressure at the crossover measurement as an adjustable parameter. There is no direct measure of the tangent heights themselves from the satellite, but rather the angles in which the instruments were pointing. Tangent heights and tangent pressures for the measurements below the crossover are determined from an iterative process.[2] This procedure is propagated downward from the crossover point (recall that the pressure at the crossover point is a parameter in the least squares fit), finding tangent height and tangent pressure for each measurement. Above the crossover point, the effects of refraction are negligible, and the tangent heights can therefore be calculated from geometry.

2.1.3 *A priori* profiles

A global fit [5] of the microwindows is carried out where the iterative process begins with the input of an initial guess or *a priori* value for each parameter. For temperature determination, *a priori* profiles from 0-150 km are derived from a combination of sources. The upper portion of the temperature profile is based on the US Naval Research Lab, Mass Spectrometer Incoherent Scatter Radar Extended Model 2000 (NRL-MSISE-00, commonly referred to as MSIS [6]) and the lower portion is based on meteorological data from the Canadian Meteorological Centre (CMC) of the Meteorological Service of Canada (MSC).

The MSIS model requires knowledge of time and spatial coordinates, as well as solar and geomagnetic indices to produce a temperature profile, which can span 0-1000 km in altitude.[6] MSIS is primarily intended for upper atmospheric work, thus does not meet our requirements for accuracy in the troposphere and lower stratosphere, leading to the need for meteorological data. The meteorological data we plan to use will be from the analysis stage of the CMC weather forecast model which involves post-processing of the forecast by assimilating measurements into the model with a delay of about 6-9 hours with respect to real-time. The meteorological data applied for each profile will correspond to the latitude, longitude and time for a tangent height of 30 km.

The technique for splicing or joining the profiles from these two different sources to ensure a smooth transition begins with determining the position of the stratopause in the MSIS profile to the nearest 1 km grid point and discarding MSIS data 3-4 km below this level. Linear interpolation is used to join the remaining MSIS modeled temperature profile to the CMC meteorological temperature profile and smoothing is applied for 2 km above and below the intersection points between each end of the interpolated data and each of the profiles, to give a single smooth profile as shown in Figure 3.

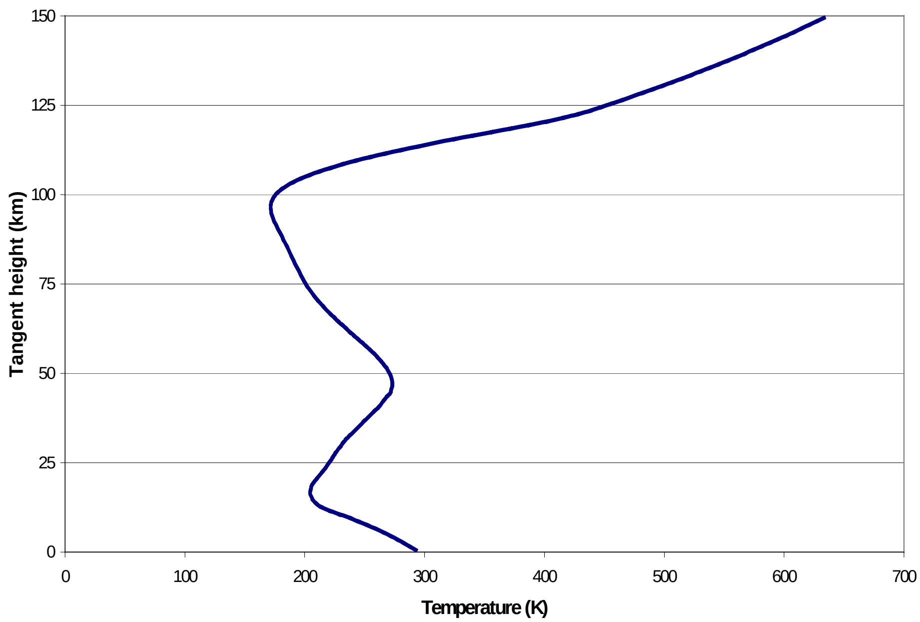

Figure 3: An example of an *a priori* temperature profile created using the method described above for 1993 April 14 at approximately 19:27 UT, -28.16º latitude and 333.18º longitude.

2.2 Volume mixing ratios (VMRs)

Pressure, temperature, and tangent heights are all established in the pressure/temperature retrieval stage. The VMR profiles of 14 baseline species and a list of other potential species will be determined. *A priori* VMR profiles for specific gases are taken from archived ATMOS [3] profiles extrapolated forward in time and appropriately adjusted for seasonal variation if necessary. Least squares fitting parameters for VMR retrievals are the VMRs at measurement tangent heights. Interpolation onto the standard 1 km grid is performed using a piecewise quadratic expression.[2] Microwindow sets are selected (where possible) to minimize spectral contamination from other species. As such, most gases are retrieved independently. For some heavier molecules, such as N_2O_5 and CFC-11, which have broad absorption features, it is not possible to find microwindows clear of interferences. Thus, a multiple molecule approach has been developed for ACE, whereby several molecules can be analyzed at the same time. The molecule of interest and the interferers are retrieved simultaneously, and then information on the interferers is discarded. The multiple molecule approach also has benefits for CO, allowing retrieval down to lower altitudes than would otherwise be possible. Ultimately, we expect to produce VMR profiles of about 30 atmospheric trace gases from the ACE-FTS measurements.

2.3 Temperature, pressure and VMR retrieval results

Development of the retrieval algorithms is nearing completion. The subroutines for creation of *a priori* temperature profiles described earlier have been tested using archived occultation points from the ATMOS missions and work well for both sunrises and sunsets at a variety of occultation times and tangent point locations (one example is shown in Figure 3). Testing of the subroutines for the temperature, pressure and VMR fits were described in an earlier work.[2] Temperature, pressure and VMR retrievals for all of the baseline molecules for the ACE mission: O_3, H_2O, CH_4, N_2O, NO_2, NO, HNO_3, N_2O_5, $ClONO_2$, HCl, CCl_2F_2, CCl_3F, HF, and CO were performed (with synthetic and ATMOS data). The results for ATMOS data were consistent with previous (ATMOS version 2 and version 3) retrieval results for the same spectra. Results for many other non-baseline molecules have also been retrieved from the ATMOS data set.

These past tests did not incorporate the *a priori* temperature profiles described above, but a complete end-to-end test of the retrieval of all parameters to be derived from the FTS spectra is expected to be carried out in the near future.

2.4 Imager Retrievals

The only quantity that can be inferred directly from imager measurements is atmospheric extinction. It is not possible to separate out contributions from different sources (as is done with the FTS measurements) because of the lack of frequency coverage. The first step in retrieving atmospheric extinction from imager measurements is to calculate transmittances from the raw measurements. As with the ACE-FTS retrievals, this is done by dividing atmospheric measurements by exoatmospheric measurements. The process is more complicated for the imagers because the calculation must be performed separately for each imager pixel, and the changing solar shape at lower tangent heights also creates challenges in mapping pixels in the atmospheric measurements to the corresponding pixels in the reference measurement: it is not a simple one-to-one mapping.

For convenience, atmospheric extinction will be retrieved on a fixed altitude grid, with grid-point spacing of either one kilometer or one-half kilometer. A non-linear least squares fitting approach will be used, determining extinction values at the grid point altitudes that best fit the set of measured transmittances. Pressure and temperature in this fitting process will be fixed at the values determined from the ACE-FTS P/T retrievals. Special consideration will be given the small set of pixels looking at the same slice of atmosphere as the ACE-FTS. In particular, these pixels on the imagers will indicate when aerosols or polar stratospheric clouds are within the FTS field of view. Contributions to the imager extinction from gas phase molecules (e.g., ozone, water, NO_2, etc.) can be calculated from the results of ACE-FTS and MAESTRO retrievals. By accounting for gas phase contributions and the contributions from Rayleigh scattering to atmospheric extinction at the imager frequencies, aerosol extinction as a function of altitude can be determined from the imager measurements.

3. INSTRUMENT CALIBRATION TESTING

The challenge in testing instrumentation destined for space-borne measurements is to characterize the performance of the instruments in a simulated space environment. In September 2002, the FTS underwent acceptance tests carried out by ABB Bomem at their facility in Quebec City. From February to March 2003, approximately six weeks were spent testing the ACE-FTS (including the imagers) and MAESTRO at the Instrument Calibration Facility (ICF) in the Department of Physics at the University of Toronto. The Toronto tests were generally intended to integrate the ACE-FTS and MAESTRO so that all instruments could be tested simultaneously under on-orbit conditions to determine their performance capability to achieve the ACE science goals. Some more specific objectives of testing were to: a) Test the performance of the FTS using the passive cooler; b) Perform gas cell measurements using the FTS and MAESTRO; c) Carry out complete testing of the solar imagers; and d) Characterize the sun-tracker pointing co-ordinates. In this work, the focus of discussion will be the Toronto tests characterizing the FTS performance, while other papers in these proceedings deal with other aspects of the testing process.

The ICF in Toronto, operated under the supervision of Prof. James R. Drummond, includes a *Class 10000* clean-room with a large thermal-vacuum chamber (TVAC). The cylindrical TVAC (shown in Figure 4) has a diameter of 2 m and is approximately 5 m long with rounded ends and an optical window on one side. During testing, the TVAC is evacuated to simulate the extremely low pressures of space. In spite of the large size of the TVAC, the vacuum system is still capable of achieving pressures on the order of 10^{-6} torr with instruments and equipment inside. During testing, the instruments were mounted to a copy of the spacecraft baseplate and surrounded by a cold shroud to simulate the low temperature environment of space. Both the baseplate and the cold shroud were mounted to a handling frame which was then mounted to the frame of the TVAC. The cold shroud is a cylindrical copper structure, open at one end. Liquid nitrogen is flowed through a series of thin copper pipes surrounding the surface of the cold shroud to cool the system. To reduce heat absorption from the surroundings, the cold shroud was wrapped in multi-layer insulation (MLI), as shown in Figure 4. (The fork assembly seen in front of the sunshield was used to rotate and place the apparatus in the TVAC and was removed prior to closing the chamber.)

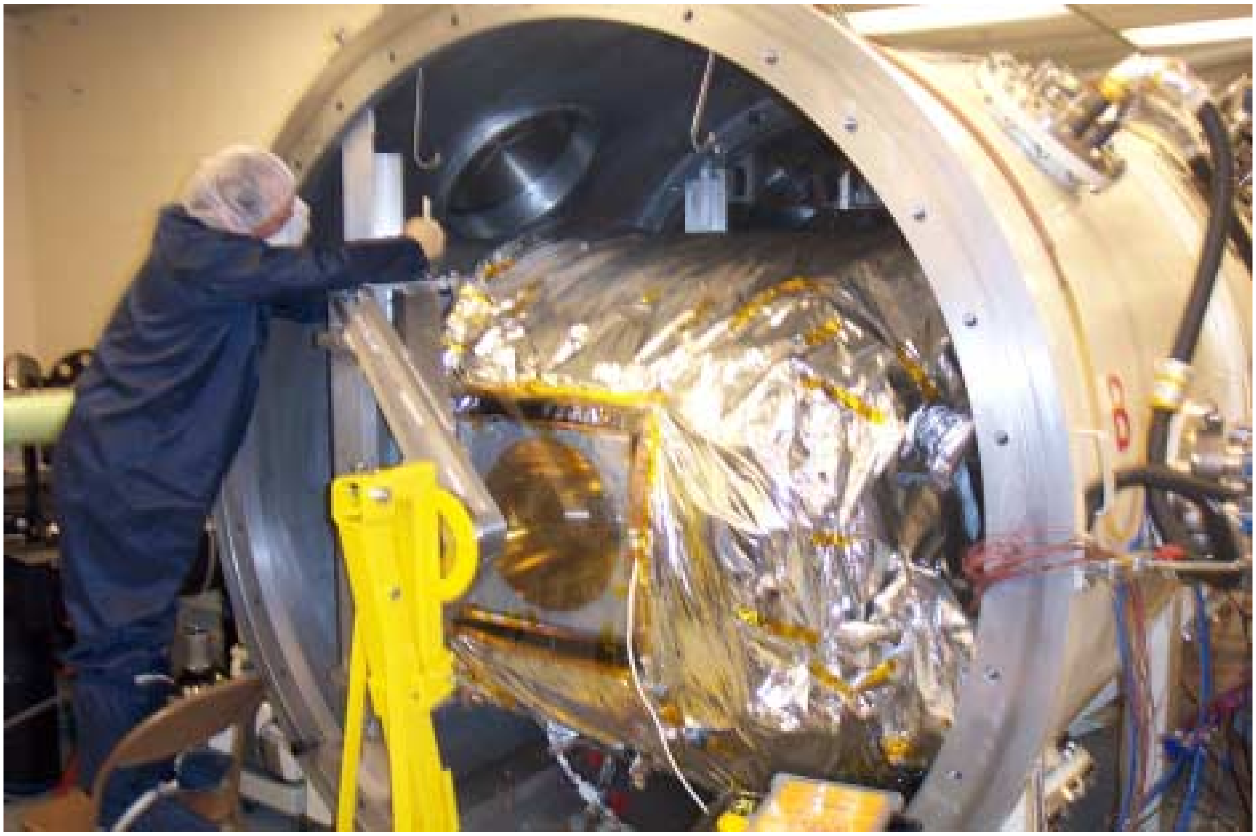

Figure 4: The instruments are enclosed in the cold shroud (wrapped in MLI) and placed in the TVAC in preparation for testing.

When mounting the apparatus in the TVAC, the optical input opening of the baseplate was aligned with the TVAC window (which cannot be seen in the above photo) and the square opening on the cold shroud was aligned with the square sunshield of the detector cooler assembly. On orbit, the passive cooler will be pointed to deep space to cool the detectors, whereas during test, the passive cooler was pointed at the cold target which simulated deep space. The cold target is a panel with a coolant reservoir behind it. In the early stages of detector cooling, the cold target was filled with liquid nitrogen, however to more closely mimic the ~3 K temperature of deep space, liquid helium (with a boiling point of 4.2 K at standard pressure) was used for further cooling. We expect the passive cooler to be able to cool the FTS detectors to less than 90 K on orbit. With liquid helium in the cold target, the detectors were cooled to 88 K and this temperature was maintained for extended periods of time while operating the FTS.

3.1 Gas cell FTS spectra

To provide radiation for both the ACE-FTS and MAESTRO, three different light sources were required. In the visible and UV, Quartz-Iodide and Mercury-Xenon lamps were used, while in the infrared, the highest temperature commercial hot blackbody (HBB) source available was employed. The HBB was capable of operating at a temperature of 3000ºC, but was limited to approximately 30 minutes at this temperature due to the lifetime of the HBB cavity. A collimator system was used to direct the radiation from the sources into the TVAC. In order to measure the absorption spectra of individual gases, a gas cell was placed in the optical path. Since no attempt was made to enclose and purge the optics to remove atmospheric impurities, strong H_2O and CO_2 features are seen, such as in the N_2O spectra shown in Figure 5, however the N_2O signal is not affected by the presence of these impurities. In the low wavenumber region of the MCT spectrum, the thickness of the baseline is not due to noise, but is a result of channelling. During the design phase of the FTS, a photoconductive MCT detector was considered, but was found to be extremely non-linear, so a photovoltaic MCT was found to replace it. This MCT has good linearity, however it is backside-illuminated and reflections from parallel surfaces within the detector chip, appear as the primary channelling in the spectra. The primary channelling has a period of ~2 cm^{-1} and a small amount of secondary channelling is also present. Fortunately, the channelling is stable for a fixed detector temperature and can be removed by ratioing sample spectra to reference spectra recorded at the same detector temperature. Examinations of narrow wavenumber ranges of our N_2O spectra also indicate the instrument lineshape (ILS) was very good with post-zero filling and prior to apodization.

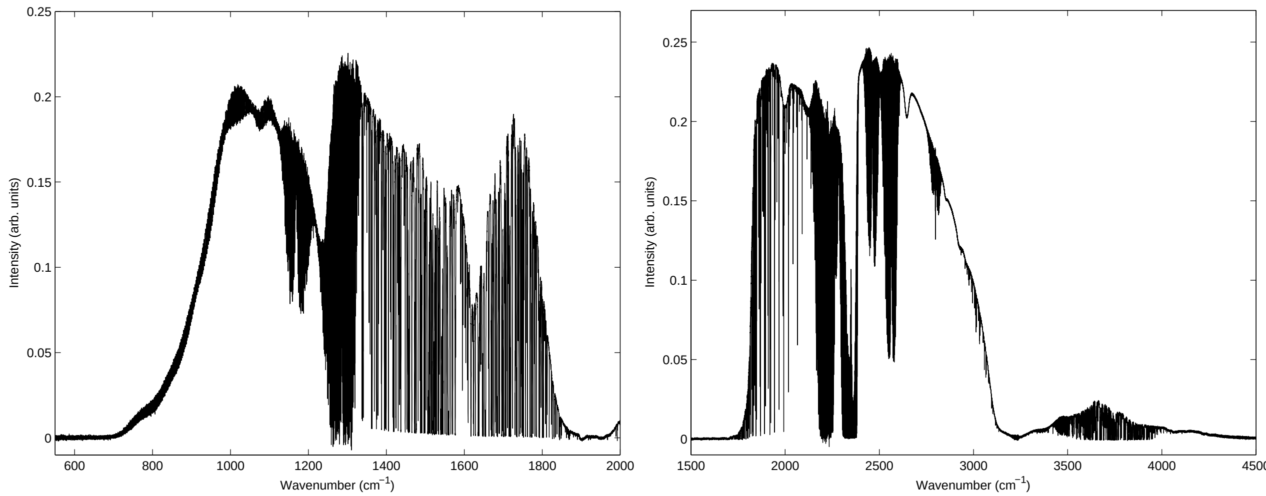

Figure 5: FTS spectra of ~12.0 torr N_2O (with typical atmospheric impurities) recorded at a resolution of 0.02 cm^{-1} in the MCT band (left) and InSb band (right).

3.2 FTS signal to noise ratio

The signal to noise ratio (SNR) requirement for the FTS is to measure CO_2 and the 14 baseline species: O_3, H_2O, CH_4, N_2O, NO_2, NO, HNO_3, N_2O_5, $ClONO_2$, HCl, CCl_2F_2, CCl_3F, HF, and CO with an SNR of 100:1, however an SNR of 200:1 was considered more desirable. Since the 3273 K HBB source used for testing purposes was not quite as hot as the photosphere of the sun at 5800 K, we must estimate the SNR which would be obtained on orbit by co-adding test spectra. Taking into account this difference in flux and the transmission of the optical setup, it was determined that ~40 co-adds were needed to simulate the SNR that will be obtained during orbit at a nominal wavenumber of 2400 cm^{-1}. Using 40 co-adds, the SNR for the MCT in the 900-1600 cm^{-1} region ranges from 200:1 to 500:1, while the SNR for the InSb in the 1800-3000 cm^{-1} ranges from 500:1 to 1200:1 (Figure 6). All species listed with the exception of HF and $ClONO_2$ have at least one microwindow, that falls in these high SNR regions. Thus for 13 of 15 molecules, the SNR requirements are easily met. The SNRs within each of the eight HF microwindows are only slightly below requirement with a range of ~30:1 to ~100:1, depending on detector temperature and the specific microwindow. Fortunately, the requirement can be met for HF since it is a stable long-lived species in the atmosphere and it should be possible to average between successive occultations to improve the effective SNR. However, for $ClONO_2$, with its single microwindow at 780 cm^{-1} near the edge of the MCT band, the SNR is only ~30:1 at a detector temperature of 89 K. Although $ClONO_2$ is stable in the atmosphere for most of the year to allow co-addition of spectra, it is variable during the polar winter and spring which are the most scientifically interesting times for $ClONO_2$ observation. It is likely that improvement could be seen for the SNR in the $ClONO_2$ microwindow if detector contamination is reduced. Furthermore, if detector temperatures below 80 K are achieved on orbit, the SNR in this microwindow would improve.

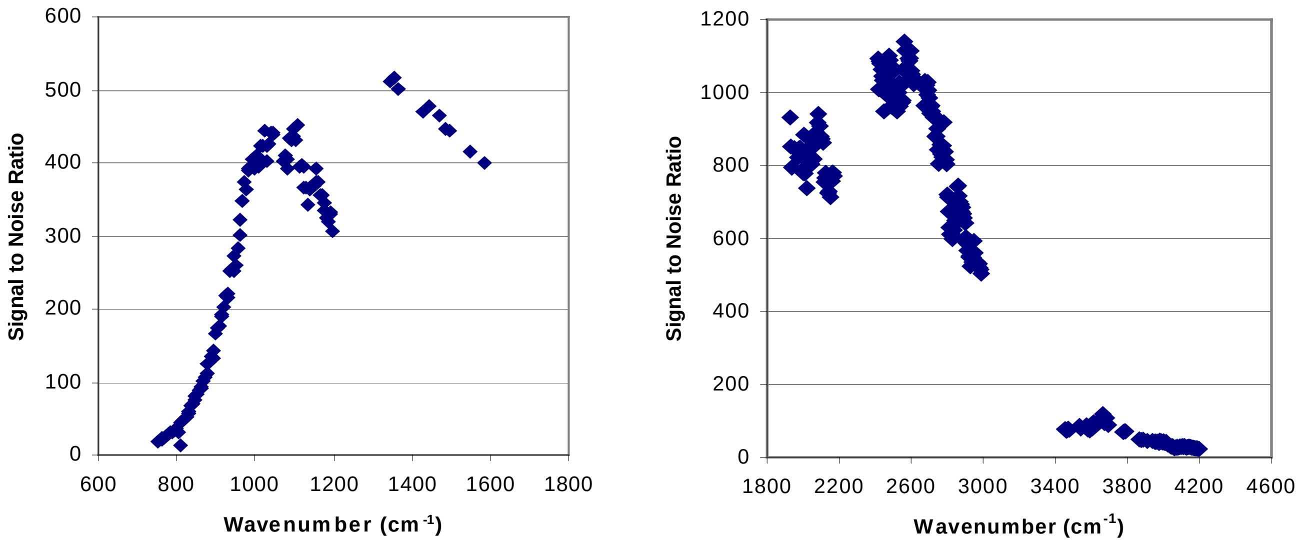

Figure 6: Plot of SNR vs. wavenumber for the MCT band (left) and InSb band (right) at a detector temperature of 89 K.

To test detector performance as a function of temperature, high and low resolution spectra with a HBB source were recorded over a period of three days as the detector temperatures decreased from 110-89 K. Selected regions of these spectra can be seen in Figure 7. Typically one would expect the performance of infrared detectors to increase as they are cooled; however, in general, it appears that the opposite is observed here. This is attributable to the build-up of contaminants on the detectors decreasing the signal transmission over time. It is most evident in the ~3000-3500 cm^{-1} range where water ice absorbs strongly, but the effects are seen to a lesser extent over the entire range of both detectors. However, there is evidence that the sensitivity of the detector itself is still improving with cooling as it should, in the ~700-790 cm^{-1} region of the MCT the improvement in sensitivity exceeds the effect of contamination (Figure 7). Similar results were obtained during the FTS acceptance test carried out by ABB Bomem and are a typical problem for space instruments with cryogenic detectors.

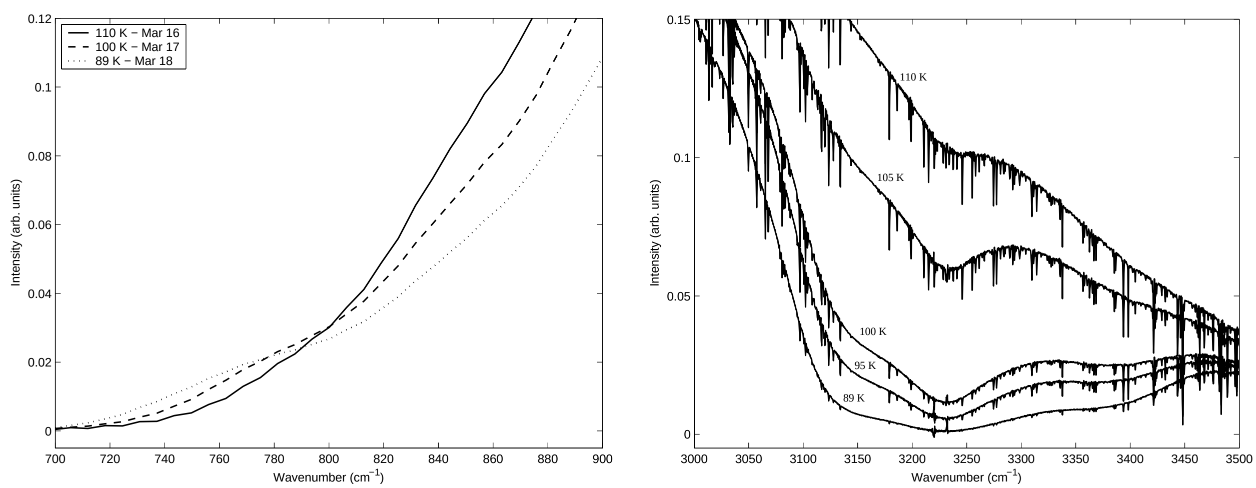

Figure 7: Portions of the spectra recorded in the MCT band (left) and InSb band (right) during detector cool-down over a period of three days. These spectra were recorded after one test decontamination cycle had been carried out. For both bands, spectra were recorded at a resolution of 0.4 cm^{-1}, however for the MCT band, the spectra are shown with the resolution degraded to 6.3 cm^{-1} to remove channelling for easier comparison. Also, spectra at two detector temperatures (95 K and 105 K) were omitted from the MCT figure for clarity, but the same trend is observed for spectra at the temperatures that have been omitted.

The method for dealing with this problem on orbit is to run decontamination cycles in which the detectors are heated in order to vaporize the contaminants which in turn diffuse out to space, followed by re-cooling of the detectors. A decontamination procedure was tested in the Toronto TVAC by heating the detectors to ~25ºC for 24 hours and then re-cooling them. This worked with temporary success, since immediately after re-cooling the detector performance improved, but the gradual build-up of contaminants resumed over time. These contaminants are mainly the result of water vapor trapped in porous materials or small cavities in the instruments and MLI.

The best way to minimize the effects of contaminants on orbit is to perform an outgassing cycle during early operations, followed by decontamination cycles as necessary throughout the mission. Thus, planning for decontamination must be incorporated into the scheduling of measurements, especially during early stages of the mission. The amount of contamination should decrease with time, thus improving SNR for all species including problem species like $ClONO_2$, but the degree of improvement on orbit can not be predicted from our tests due to many factors such as the amount of initial contaminants present on orbit and changes in MLI-wrapping of spacecraft and instrument components.

4. THE ACE DATABASE AND DATA STORAGE HARDWARE

During the nominal two-year ACE mission, there will be as many as 20,000 measurement opportunities. Several terabytes of measured data will be produced, and it will need to be stored, processed and then made available for distribution to the mission's international science team. The approach used for the ACE database is that of a central high-throughput server which is supplemented by less expensive processing nodes and storage.

Measured data is collected from the spacecraft at one of two Canadian Space Agency facilities: St. Hubert, Quebec and Saskatoon, Saskatchewan. The Mission Operations Center (MOC, also located in St. Hubert), processes all downlinked data and makes the science payload files available for transfer to the science team.

The Science Operations Center (SOC), located at the University of Waterloo, houses the mission's data processing server and archive. The operations server is an enterprise-class Sun Fire 3800, with four processors and 4 GB of memory. The server's extremely high throughput and availability allow it to serve as the "backbone" of the entire ACE processing system.

Main storage for the mission database is a Sun StorEdge T3 array with a capacity of 500 GB. This high-performance array uses a Fiber Channel bus to achieve simultaneous read and write operations at 40 MB/s. A separate 1.2 TB PC-based server provides archive storage, with a reduced bandwidth of approximately 5 MB/s. One can maximize the total storage of the system by segmenting the database (below).

The system's backup needs are met with a 40-cartridge tape library, which includes two DLT8000 drives and a robotic interchange for autonomous operation. The library has an internal capacity of 1.6 TB, but the backup software (Legato Networker) can manage off-line volumes so the backup process can be extended without limit.

The operations server and storage are housed in a single high-availability rack with redundant power busses and ventilation.

4.1 Database Software

The chosen database software is the PostgreSQL relational database management system. The server backend software has been compiled from source and tuned to provide maximum performance on the ACE hardware. As mentioned above, the database can be segmented to match the performance of the storage hardware:

- Internal Data, which are the tables and indices used by the backend server, are stored on the high-performance array to maximize the response of the database itself.

- Current data that is undergoing processing, also resides on the main storage array. The entire set of data products, which is much smaller than the corresponding measurements, is given high priority so that it's available for trending or averaged analysis.
- Housekeeping data from the spacecraft bus is stored verbatim in the high-performance array. This will make it possible to correlate instrument performance with any of the telemetry points that are monitored by the spacecraft. For example, it might be necessary to analyze thermal effects while the instruments were operating.
- Archive data, which are datasets that have already been processed, are moved to the secondary array. The time period after which data will be considered "archive" will be a function of available space on the main array, but should be at least a few months. Access to this relatively slow array is still orders of magnitude faster than on an offline medium such as tape.

PostgreSQL was originally chosen for its open-source license and space heritage, but development of software for the ACE mission has been bolstered by the variety of interfaces to the PostgreSQL system. Core data processing software written in C and FORTRAN use the database instead of traditional input/output files, and higher-level distribution can be performed using Java or by the web server. By using a consistent storage approach throughout data processing, it is much simpler to tie all of the software components together for end-to-end distribution.

Connections to the database server are made using TCP/IP network sockets, so remote access to the data is the same as if the application was running on a local server. Username and password pairs are used to authenticate each connection, and additional measures are taken at the network level to prevent unauthorized access.

The database's wide variety of built-in data types is also of great help when designing a database as large and complex as that of an entire scientific mission. Of particular use are Binary Large Objects (BLObs), which have the same properties of normal UNIX files, but are stored within the database. Multi-megabyte objects such as interferograms and spectra are cast as BLObs in the database, and the array datatype is used for smaller structures.

4.2 Data Processing

In addition to the operations server, a small cluster of compute nodes has been installed to provide additional processing. The current cluster employs four Sun Fire V100 servers, each with 1 GB of memory, but practically any UNIX system can be added at a later date. To prevent interference with other server traffic, a dedicated 100 Mbps network connects the cluster to the operational database.

Job queuing is managed with the Portable Batch System (OpenPBS), which greatly facilitates the management of the entire system. As data sets arrive, processing jobs are submitted to the queuing system, which then distributes the execution sequence amongst the available nodes. It is worth noting that a database approach greatly simplifies such a distributed processing system. As mentioned above, database transactions are inherently networked, and so no special file handling is required to process data on separate hosts.

The principal benefit of a cluster approach is that the compute power of a new node contributes directly to the total processing, with linear cost and minimal administration. Additionally, this approach is much more modular than having a single multiprocessor system. Once the mission is underway, the scheduling software does not differentiate between different operating systems, so processing can be shared across a cluster of mixed hardware. Work is underway to verify the performance of data processing components on Intel-based Linux systems. Reliability is not a real concern within a cluster, so one can use low-end personal computers to improve the performance of the entire system.

5. CONCLUSIONS

We have summarized the process for retrieval of temperature and pressure from ACE-FTS spectra based on two possible scenarios: good pointing knowledge and poor pointing knowledge. A technique for creation of automatic *a priori* temperature profiles by combining meteorological data with modeled data has been described, and VMR and imager retrievals have also been described briefly. Our FTS retrieval tests of temperature, pressure and VMRs thus far indicate promising results for all areas with a full end-to-end test of the entire forward model planned for the near future.

Instrument testing has shown that passive cooling of the FTS detectors functions as intended and that the FTS is capable of recording good quality spectra. Although the MCT detector exhibits channelling, it appears to pose no significant problem if dealt with by ratioing to the appropriate reference spectra. Spectra in both the InSb and MCT bands had excellent SNRs over nearly the entire desired spectral range; however two molecules (HF and $ClONO_2$) of the 14 baseline species may have some complications with meeting the SNR requirement of 100:1. Averaging between occultations will solve the problem for HF but is only a partial solution with $ClONO_2$ because its VMR is not stable enough for averaging during the most scientifically interesting times. Fortunately, as the mission proceeds, SNR should likely improve upon that estimated during testing by running repeated decontamination cycles, but the degree of improvement is difficult to predict.

When ACE is in orbit, the mission will generate large amounts of scientific data that will be handled by an enterprise-class Sun Fire 3800 acting as the operations server. Main storage for the mission database is a Sun StorEdge T3 array with a capacity of 500 GB. For backup storage, a tape library with a 1.6 TB internal capacity will be used with the potential for managing offline volumes, so storage capacity may be extended without limit. The chosen database software is the PostgreSQL relational database management system. A small cluster of compute nodes has also been installed for additional processing. With this modular approach provided by a cluster, the most current technology can be added to the system to suit our potentially changing needs at later stages of the mission.

ACKNOWLEDGEMENTS

Primary funding for the SCISAT-1/ACE mission was provided by the Canadian Space Agency. Support was also provided by the Natural Sciences and Engineering Research Council of Canada (NSERC) through a collaborative research grant and an industrial research chair. We would like to thank all the members of the ACE Calibration Test Team and the University of Toronto Instrument Calibration Facility for their assistance in ACE testing, especially Prof. James Drummond and Denis Dufour. We would also like to give credit to the ACE Project Manager, Dr. Mike Butler for his contribution to ACE testing and many other aspects of the mission.

REFERENCES

1. Rodgers, C.D. *Inverse Methods for Atmospheric Sounding,* World Scientific, New Jersey, 2000.

2. Boone, C. & Bernath, P. “Pressure/temperature and volume mixing ratio retrievals for the Atmospheric Chemistry Experiment” *Proceedings of SPIE, Earth Observing Systems VII*, Editor W.L. Barnes, vol. 4814, pp. 50-61, 2002.

3. Gunson, M.R., Abbas M.M., Abrams M.C., Allen M., Brown L.R., Brown T.L., Chang A.Y., Goldman A., Irion F.W., Lowes L.L., Mahieu E., Manney G.L., Michelsen H.A., Newchurch M.J., Rinsland C.P., Salawitch R.J., Stiller G.P., Toon G.C., Yung Y.L. & Zander R. "The Atmospheric Trace Molecule Spectroscopy (ATMOS) experiment: Deployment on the ATLAS Space Shuttle missions" Geophysical Research Letters, **23**, 17, pp. 2333-2336, 1996.

4. Roble, R.G. “On the feasibility of developing a global atmospheric model extending from the ground to the exosphere”, Geophysical Monograph 123, American Geophysical Union, 2000.

5. Carlotti, M. “Global-fit approach to the analysis of limb-scanning atmospheric measurements” Appl. Opt. **27**, 15, pp. 3250-3254, 1988.

6. Hedin, A.E. “Extension of the MSIS Thermosphere Model into the Middle and Lower Atmosphere, J. Geophys. Res., **96**, A2, pp. 1159-1172, 1991. (http://uap-www.nrl.navy.mil/models_web/msis/msis_home.htm)