



SAGE version 7.0 algorithm: application to SAGE II

R. P. Damadeo¹, J. M. Zawodny¹, L. W. Thomason¹, and N. Iyer²

¹NASA Langley Research Center, Hampton, VA, USA

²Science Systems and Applications Inc., Hampton, VA, USA

Correspondence to: R. P. Damadeo (robert.damadeo@nasa.gov)

Received: 16 May 2013 – Published in Atmos. Meas. Tech. Discuss.: 10 June 2013

Revised: 7 November 2013 – Accepted: 8 November 2013 – Published: 17 December 2013

Abstract. This paper details the SAGE (Stratospheric Aerosol and Gas Experiment) version 7.0 algorithm and how it is applied to SAGE II. Changes made between the previous (v6.2) and current (v7.0) versions are described and their impacts on the data products explained for both coincident event comparisons and time-series analysis. Users of the data will notice a general improvement in all of the SAGE II data products, which are now in better agreement with more modern data sets (e.g., SAGE III) and more robust for use with trend studies.

1 Introduction

The Stratospheric Aerosol and Gas Experiments (SAGE I, II, III/METEOR-3M, and III/ISS) are an ongoing series of satellite-based solar occultation instruments spanning over twenty-six years. Measurements from the SAGE series have been a cornerstone in studies of stratospheric change, including having played a key role in numerous international assessments (e.g., WMO, 2011). Given the importance of the data, it is imperative that the data sets, and the processing codes that produce them, be maintained and, when necessary, updated and improved to reflect the evolving “best practices” for processing occultation data to science products. To facilitate using data from multiple instruments to investigate long-term variability in atmospheric components, it is important to maintain consistency in methodology (when applicable) and fundamental assumptions made in processing data from each instrument. This paper describes the first standard algorithm to process SAGE data, SAGE version 7 (v7.0). The basis of the version 7.0 algorithm derives primarily from the SAGE III/M3M version 4.0 algorithm and is intended to form the basis for the reprocessing of all members of the

SAGE series, including the Stratospheric Aerosol Measurement (SAM II) instrument. This paper provides an overview of the instrument operation and algorithm, followed by a detailed description of each step of the processing algorithm and how it is applied to SAGE II, including differences between the previous (v6.2) and current (v7.0) versions.

1.1 Instrument operation

SAGE II operated on board the Earth Radiation Budget Satellite (ERBS) from its launch in October 1984 until its retirement in August 2005. It employed the solar occultation technique to measure multi-wavelength slant-path atmospheric transmission profiles at seven channels during each sunrise and sunset encountered by the spacecraft. The optical properties of most channels were defined by the position of exit slits along a Rowland spectrometer where photodiodes measured the impinging light. The seven channels, in channel number order, were nominally located at 1020, 935, 600, 525, 452, 448, and 386 nm. Due to limitations on the size of the diodes (i.e., placing them next to each other), channels 2 and 5 were placed at the zero order location with filters providing the desired band-pass. Channel 6 required a narrow band-pass and also employed a filter to relax the requirements on high tolerance mechanical positioning of the channel 6 exit slit.

The SAGE II instrument was oriented towards nadir on the spacecraft such that the optics could observe Earth’s limb. Prior to the expected start of an occultation (event), the scan head/telescope/spectrometer assembly rotated towards the predicted azimuth location of the Sun and then locked onto the solar centroid brightness. An elevation scan-mirror then began moving the field-of-view (0.5 by 2.5 arc-minutes) across the solar disk normal to the Earth’s surface. As the field-of-view went off the edge of the Sun, the scan-mirror would reverse direction. In this way, the

field-of-view scanned vertically across the Sun while each channel recorded solar irradiance data (count values) at a rate of 64 Hz (packets per second) to construct a series of solar limb-darkening curves (counts observed as a function of time). The instrument continued this process until the Sun disappeared below Earth's limb (for sunsets) or a pre-set amount of time had elapsed (for sunrises). The benefit of scanning back and forth across the Sun, as opposed to simply staring at the Sun, is that during the course of a sunrise or sunset, the instrument was able to observe the same altitude multiple times, albeit through slightly different viewing geometries. Stated differently, the instrument was able to observe the same point on the solar disk through multiple ray tangent altitudes. This allowed for excellent vertical resolution (~ 1 km) in retrieved SAGE II data products. (Mauldin et al., 1985)

1.2 Algorithm overview

The general approach of the version 7.0 algorithm is the same as previous versions. The process begins with the production of slant-path transmission profiles at each wavelength, followed by the separation of the spectral information into species slant-path column abundances, and finally, the inversion to species density and aerosol extinction profiles. SAGE II processing begins with the assimilation of instrument data (time of each packet, scan-mirror elevation position of each packet, and count value in each channel of each packet), spacecraft and solar ephemeris data, meteorological data for the time and location of the observation, and spectral information related to trace gas species involved in the retrieval process.

Spacecraft and solar ephemeris data are processed to determine where the spacecraft was, where the Sun was, and what the viewing geometry was at all times during the event. Meteorological data are processed to create vertical profiles of temperature and density that are used to calculate refraction angles that are subsequently used to construct the proper refracted viewing geometry observed by the instrument. As an additional preprocessing step to facilitate later calculations, spectral characteristics of the Sun, atmospheric molecular scattering, and trace gas species are combined with the spectral filtering characteristics of each channel to create band-pass averaged effective cross-sections observed by each channel.

The solar limb-darkening curves are used to determine the location of the physical edges of the Sun, which, when combined with ephemeris data, allow for an accurate mapping of instrument data to viewing geometry data. The instrument observed every point on the Sun both through the attenuated atmosphere and high above it, which allowed for SAGE II to be self-calibrating. Exoatmospheric scans are combined to create an exoatmospheric limb-darkening curve (or I zero curve). Scans observed through the atmosphere are then compared to the standard I zero curve; each point on the surface

of the Sun as viewed through the atmosphere is calculated as a ratio to the same point in the I zero curve to create slant-path transmission data for each channel. Several calibrations and corrections are included in the iterative processing of transmission to compensate for biases in the edge time calculations, reflectivity of the scan-mirror as a function of elevation position, rotation of the field-of-view on the solar disk due to orbital motions, and other minor corrections. Transmission data are then interpolated to a 0.5 km grid to facilitate the inversion process.

Slant-path transmission profiles are converted into optical depths and then combined with spectroscopy data to separate into species-specific slant-path optical depth profiles. These are then inverted to produce vertical profiles of ozone (O_3), nitrogen dioxide (NO_2), water vapor (H_2O), and aerosol extinction (at 1020, 525, 452, and 386 nm) using a simple onion-peeling technique. In order to do this, the inversion algorithm must make the assumption that the layer of atmosphere at each altitude is homogeneous, or at least has a constant gradient, through the whole swath that the instrument observes. This assumption has obvious limitations in the troposphere and well-understood biases at higher altitudes for certain species due to rapid photochemistry, giving rise to a nonlinear variation across the terminator (Chu and Cunnold, 1994), but works well through most of the stratosphere (Cunnold et al., 1989). A general outline of the algorithm can be seen in Fig. 1. The sections that follow provide greater detail on the various steps just outlined and note when and how these differ from the approach of SAGE II version 6.2.

2 Preprocessing

Before anything can be done with the solar irradiance data measured by the instrument, a series of preprocessing steps must be performed so that this data can be placed in the proper context. The algorithm needs to determine where the instrument was looking during the collection of each data packet. This requires processing spacecraft orbital ephemeris data, along with atmospheric refraction information and the requisite meteorological information. In addition, some calculations related to the spectroscopy of the instrument are performed to facilitate later calculations.

2.1 Ephemeris

NASA's Tracking and Data Relay Satellite System measured the state vectors (position and velocity) of the ERBS on a regular basis. The Operation Support Computing Division at the Goddard Space Flight Center assimilated this data with an accurate model to determine the spacecraft position and velocity at 60 s intervals, which was provided as level zero ephemeris data to the SAGE processing team. From this original ephemeris data, events with beta angles (defined as declination of the Sun as measured from the orbital plane) greater

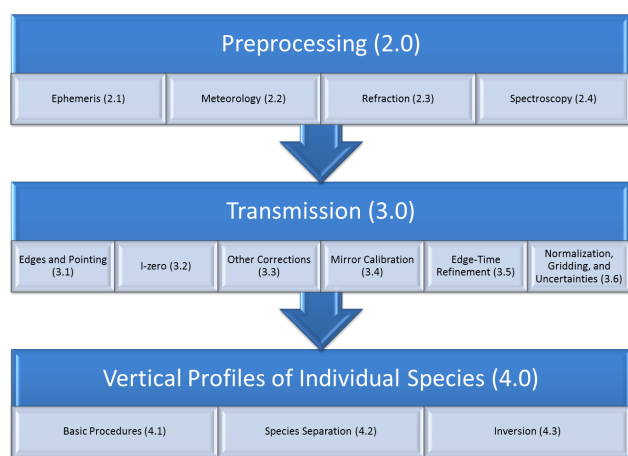


Fig. 1. A general outline of the SAGE II algorithm. The numbers correspond to sections in this paper.

than 61° or events that occur during spacecraft viewed solar eclipses are excluded from further processing. High beta angles are excluded because the duration of an event (and subsequent ground track) becomes long and the assumption of atmospheric spherical symmetry, used for the inversion process, breaks down. For each remaining event, the original spacecraft state vectors just prior to the start of the event are used as input to the processing algorithm. Using a model for Earth's gravitational potential, the equations of motion are calculated and the state vectors are propagated throughout the event. In addition to the state vectors, the position of the Sun is calculated for each time. Then, for each time, coordinate and geometrical data required by the algorithm are calculated. This includes information related to the spacecraft (sub-spacecraft latitude, longitude, and altitude), the Sun (angular size, right-ascension, declination, and azimuth), and the tangent point (altitude, latitude, longitude) looking at the center, top, and bottom of the Sun.

The methodologies for these calculations originate from Buglia (1988). For the most part, these methods are straightforward geometry but, in pre-version 7.0 algorithms, some aspects (e.g., sidereal time, precession, nutation, and Sun position) were approximated using numerical constants derived from pre-1984 data. It is important to note that in 1984, national and international almanacs adopted a revised set of physical constants put forth by the International Astronomical Union in 1976, which necessitated a change to various numerical constants used in these approximations. However, these changes were not uniformly implemented and data quality in versions prior to version 7.0 was adversely impacted by an inconsistency in ephemeris epoch usage.

In version 7.0, all physical constants (e.g., Earth's equatorial and polar radii and gravitational constant) have been updated to those used in the World Geodetic System 1984 (WGS84, updated in 2004) (NIMA Technical Report, 1997), which is the current standard. The constants used for the

multi-pole expansion of Earth's gravitational potential have been updated to those from the Earth Gravity Model 1996 (EGM96) (Lemoine et al., 1998). In addition, coded routines that relied upon approximations based upon pre-1984 data have been modified to utilize SDP Toolkit routines (Noerdlinger, 1995). Most of these changes result in inconsequential changes to the results with one notable exception. The original routine designed to calculate the position of the Sun suffered from inconsistent ephemeris epoch usage and outdated numerical approximations. The adoption of a Toolkit routine has corrected what was a previously unknown quasi-random error in the altitude registration in the SAGE II data products. For any given event, this correction manifests itself as an altitude offset between SAGE II versions 6.2 and 7.0. This altitude offset can be positive or negative, can have a magnitude up to a few hundred meters, and varies from event to event. While there is no simple dependence upon beta angle or latitude, there is some correlation with time of year, as expected, as shown in Fig. 2.

2.2 Meteorology

SAGE processing algorithms require ancillary meteorological data that relates density and temperature to altitude. SAGE II version 6.2 used multiple sources of data to yield density and temperature data (i.e., P/T data) from the surface up to 100 km. NCEP reanalysis data (Kalnay et al., 1996) was used first, yielding P/T data from 1000 mbar up to 10 mbar (~ 30 km). Above that, operational model data provided by NCEP was used up to 0.4 mbar (~ 50 km). Lastly, the Global Reference Atmospheric Model-1995 (GRAM-95) (Johnson et al., 1995) was used to extend the profiles up to 100 km. Since each atmospheric layer is assumed to be uniform, a single location was chosen to retrieve P/T data from, namely, the 20 km sub-tangent latitude and longitude. The temperature and pressure data sets were then combined and interpolated to a standard 0.5 km-spaced altitude grid.

Analysis of the version 6.2 P/T data revealed a few anomalies where NCEP operational model data were (and are) absent, resulting in the use of GRAM-95 data deep into the stratosphere. In addition, an analysis comparing SAGE II and III coincident events revealed that SAGE II NCEP operational model temperature data were typically warmer than SAGE III, even though both were using NCEP data sources (Fig. 3). The cause of this bias is unexplained. To maintain as much consistency as possible, a long-term self-consistent meteorological data set was required that spans the lifetimes of all of the SAGE instruments and provides data throughout the stratosphere. The Modern Era-Retrospective Analysis for Research Applications (MERRA) (Rienecker et al., 2011), based on the Goddard Earth Observing System Model (GEOS-5.2) (Rienecker et al., 2008), provides temperature data at 42 pressure levels from the surface up to 0.1 mbar (~ 65 km) from 1978 to the present. While this has not yet been implemented in SAGE III, there are plans to use




Fig. 2. Tangent height registration differences between version 7.0 and 6.2 for the same time index near a tangent height of 60 km for sunrises/sunsets (top/bottom) as a function of time of year. A single point is plotted for each SAGE II event.

MERRA data for the reprocessing of SAGE III/M3M data to the version 7.0 standard and to use operational GEOS data for the upcoming SAGE III/ISS mission. SAGE II version 7.0 uses MERRA data from the surface up to 0.1 mbar and GRAM-95 above that up to 100 km. Instead of using the actual GRAM-95 temperature values, the GRAM-95 lapse rate is used to extrapolate above the MERRA data. This is done because the MERRA lower mesospheric temperature values are often smaller than those from GRAM-95. Any attempt to merge the two data sets can introduce an artificial inversion layer in the lower mesosphere.

In both versions, after P/T profiles are determined, number density profiles are calculated. In order to have uncertainty estimates in derived quantities in the inversion process (e.g., uncertainty estimates for Rayleigh scattering slant-path extinction), we calculate an uncertainty estimate in this density calculation, based on the uncertainty in the original temperature data. NCEP reanalysis does not provide uncertainty estimates for their model products; however, we obtained some temperature uncertainty values (ranging from 4 K in the troposphere to 15 K in the mesosphere) to use at each pressure level (Finger et al., 1993). MERRA also does not provide uncertainty estimates for their model products. We continue to discuss the topic of model uncertainties with the MERRA group. In the meantime, we are currently adapting the NCEP uncertainty values as a rough guideline.

2.3 Refraction

Our refraction algorithm calculates the elevation angle (relative to the local horizontal plane at the position of the spacecraft) of the refracted Sun (where the instrument sees the Sun) and the total refraction angle of the light ray as a




Fig. 3. Coincident event temperature profile comparisons between SAGE III and SAGE II. The biases in the middle stratosphere between SAGE III and SAGE II v6.2 data are unexplained. The comparison with v7.0 illustrates differences in the shape of the upper stratosphere between MERRA and NCEP/GRAM-95, with MERRA typically having a less well-defined transition to the mesosphere and a cooler and/or higher (in altitude) stratopause.

function of wavelength and possible tangent heights assuming spherical geometry (Fig. 4). It also calculates the layer slant-path matrix with which it determines the total number density of the slant-path air column (mass path) by integrating along the curved path of the light ray (again with a spherical model). The use of a spherical model for these calculations provides a good first guess before refinement using an oblate Earth model can be done. These parameters are chosen because, for each packet, the algorithm needs to know precisely where on the Sun the instrument is pointing and these parameters allow it to easily go between where the instrument is pointing in space and where the instrument is looking on the face of the Sun. The methodology of computation remains largely unchanged from version 6.2 and comes from Chu (1983) and Auer and Standish (2000). After refraction, tangent point altitudes, latitudes, and longitudes are updated, taking an oblate Earth model into account.

There is one important error in previous versions to note, which is corrected in version 7.0. As the algorithm determines the refraction angle, it begins at the surface and moves to higher tangent altitudes. At a small threshold value for the refraction angle, this process ceases and remaining refraction angle values are assumed to be zero. However, these values were not explicitly set to zero. The result of this was that if one event hit this threshold at some altitude and the following event hit this threshold value at some lower altitude, then the refraction angle values between these two altitudes would be carried over from the previous event, potentially through multiple events. This generally occurred in the 50–60 km




Fig. 4. SAGE refracted viewing geometry.

range, and would be interpreted as non-physical characteristics in the atmosphere, which could propagate downward in the inversion process. However, while this potentially introduced random deviations in derived products (e.g., as large as a few percent in ozone), the effect averaged out, even within a single event, to zero. As such, correcting this error in version 7.0 did not uncover or correct any biases in prior versions.

2.4 Spectroscopy

The retrieval process requires absorption cross-sections for O_3 , NO_2 , the oxygen dimer (O_2-O_2), and molecular (Rayleigh) scattering at all wavelengths. This is done by first combining each channel's measured spectral response function with the solar spectrum (Kurucz et al., 1984). The spectral response of channels 2, 5, and 6 also incorporate the long-term evolution in the spectral characteristics of their respective band-pass filters. These spectral responses are then combined with cross-section data for each species (i.e., O_3 , NO_2 , O_2-O_2 , and Rayleigh) and integrated across each channel's spectral range. Rayleigh cross-sections ($cm^2 molecule^{-1}$) are wavelength-dependent and are calculated in the same fashion as in Bucholtz (1995). O_2-O_2 cross-sections ($cm^5 molecule^{-2}$) are taken to be wavelength-dependent, are assumed to scale with density, and are calculated in the same fashion as Mlawer et al. (1998) (for wavelengths encompassing channels 1 and 2) and Newnham and Ballard (1998) (for wavelengths encompassing channels 3 through 7). O_3 and NO_2 cross-sections ($cm^2 molecule^{-1}$) are taken to be both wavelength-dependent and temperature-dependent. In version 6.2, they were derived from the Shettle and Anderson cross-section compilation (Shettle and Anderson, 1995), which is the same cross-section compilation that was used in SAGE III version 3.0. In SAGE III version 4.0 (Thomason et al., 2010), however, the O_3 cross-sections were updated to the Bogumil (Scanning Imaging Absorption Spectrometer for Atmospheric Chartography (SCIAMACHY) V3) cross-sections (Bogumil et al., 2003), which had a positive impact on several data products

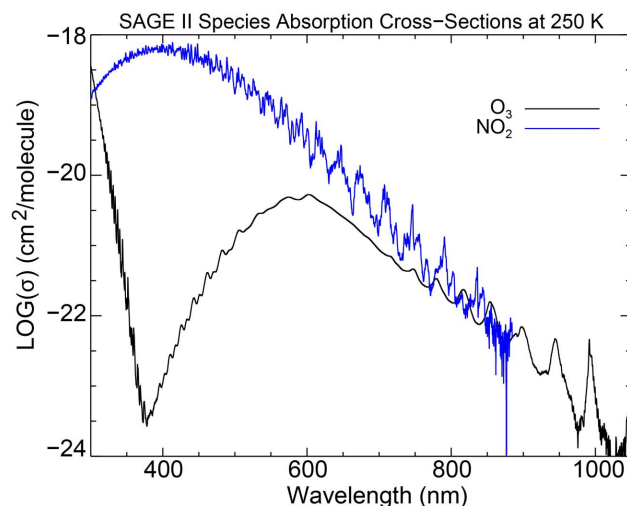


Fig. 5. Absorption cross-sections for SAGE II retrieved species at 250 K over SAGE II measurement wavelengths for v7.0.

and produced better agreement with in-atmosphere measurements (Pitts et al., 2006). As such, SAGE II version 7.0 uses the SCIAMACHY V3 O_3 cross-sections, including companion measurements of NO_2 absorption cross-sections encompassing SAGE II measurement wavelengths (Fig. 5). These band-pass averaged, or effective, cross-sections are combined with vertical density (for O_2-O_2) or temperature profiles (for O_3 and NO_2) to create vertical “profiles” of effective cross-sections for these species in each channel, facilitating later computations.

Effective cross-sections are constructed in this fashion because the interactions between species absorption within the band-pass of any channel are negligible. The O_3 retrieval is primarily dependent upon observations in channel 3 (600 nm), where the NO_2 cross-section has decreased two orders of magnitude from its peak and contributes a negligible amount to the extinction. The NO_2 retrieval is primarily dependent upon correlative measurements in channels 5 (452 nm) and 6 (448 nm), where the O_3 cross-section has decreased nearly two orders of magnitude from the Chappius band and displays little to no structure within these narrow band-passes. Aerosol retrievals have a strong dependence upon observations in channel 1 (1020 nm), where the cross-sections of both O_3 and NO_2 have decreased by several orders of magnitude and contribute a negligible amount to the extinction. The use of effective cross-sections requires a lack of spectral correlation between species within a given channel. The reason for this is that, at large optical depths, absorption features at spectral scales finer than the instrument can resolve will become saturated, and the process of removing a single species' contribution to a channel becomes nonlinear.

Due to the nature of the water vapor feature near 940 nm, the creation of effective cross-sections for water vapor is not possible and the absorption must be modeled. Water vapor

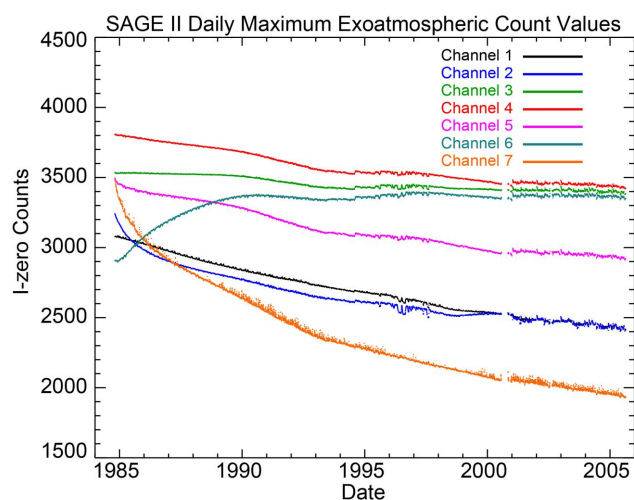


Fig. 6. Daily maximum exoatmospheric “*I* zero” count values for each channel. While a general degradation in response can be seen in each channel over time, a few channels, namely 2, 6, and 7, went through periods of rapid change during the beginning of the mission.

line data for version 6.2 were provided by L. Brown (personal communication, 2002), which were later incorporated into the water vapor line data for the 2004 version of the HITRAN (high resolution transmission) molecular spectroscopic database (Rothman et al., 2004). The version 7.0 water vapor line data come from the 2008 version of HITRAN (including the 2009 update to water vapor) (Rothman et al., 2009). These line data are used to precompute derivatives of absorption as a function of temperature, pressure, and line-of-sight molecular number density for later use with an emissivity curve-of-growth approximation (EGA) as the forward model for water vapor absorption (Gordley and Russell, 1980).

After SAGE II had been operating for some time, a long-term analysis of the water vapor product showed it to be in poor agreement with other satellites and ground measurements. While many possibilities were considered, and eventually ruled out, it was determined that the poor quality of water vapor data was a result of a shift in the spectral response of the water vapor channel (channel 2) prior to 1986 (Fig. 6). The primary reason behind this thinking was an incident associated with the SAGE II instrument during one of its final thermal vacuum tests just prior to being shipped off for integration with the ERBS. An incomplete shutdown of the cooling system caused condensation within the instrument. It is believed that the channel filters absorbed water and then subsequently dried out on orbit, affecting their spectral characteristics. Unfortunately, it was impossible to reproduce this in the lab as no original filter material remained and the company that created it was no longer in business. Version 6.2 was the first version to attempt to account for an apparent shift in the location of the water vapor channel filter. In order

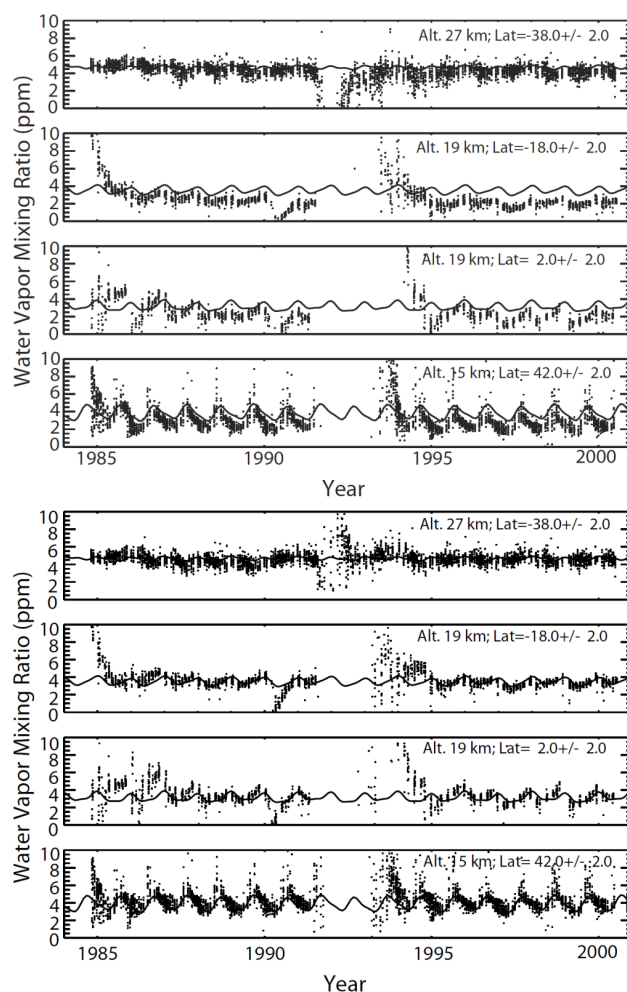


Fig. 7. SAGE II v6.1/v6.2 (top/bottom) water vapor data (dots) and HALOE climatological data (lines) (Thomason et al., 2004).

to attempt to model the filter characteristics, it was decided to adjust the spectral response of the water vapor channel to make the mean SAGE II water vapor data at one latitude and season agree with a HALOE (UARS (Upper Atmosphere Research Satellite) Halogen Occultation Experiment) (Russell et al., 1993) climatological profile for the same latitude and season (northern mid-latitudes in March). The center wavelength and full-width at half-max (FWHM) were adjusted and the two data sets were again compared. While it was impossible to completely match these data sets, it was found that the best match came from a center wavelength shift of +10 nm (945 nm) and an increase in the FWHM of 10 % (22 nm). This was then applied to all SAGE II data from 1986 onward. The before and after comparisons of that work can be seen in Fig. 7. (Thomason et al., 2004)

With the adoption of the SCIAMACHY V3 ozone cross-section database came a large shift in the Wulf ozone bands that span the SAGE II water vapor channel spectral response (Fig. 8). Given the sensitivity of the water vapor retrieval

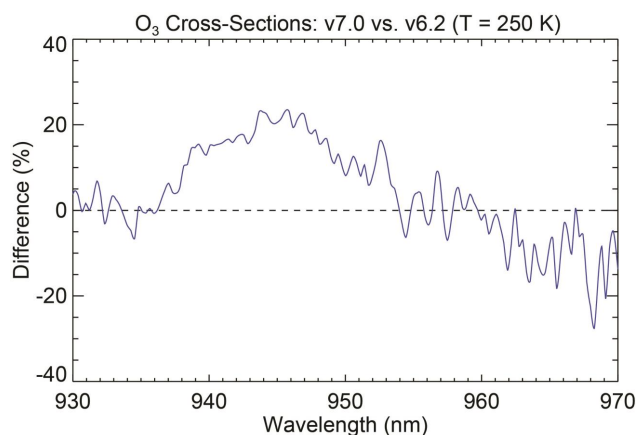


Fig. 8. SCIAMACHY V3 O₃ cross-sections minus the Shettle and Anderson compilation at 250 K in the range of 930–970 nm.

to the details of ozone (Chu et al., 1993), the evolution of the water vapor channel spectral response was reevaluated. A comparison of SAGE II version 6.2 with newer data sets (Fig. 9) shows that, while version 6.2 matches well to HALOE, as expected, there is an offset of about 10 % with both SAGE III and the MLS (Microwave Limb Sounder) (Lambert et al., 2007) on board the Aura satellite, perhaps suggesting the choice of HALOE as a standard for inferring channel drift was not optimal. The channel 2 drift assessment was repeated, except SAGE III water vapor was used to infer the location of the water vapor channel rather than HALOE. The smallest difference between data sets came from an additional center wavelength shift of +2.7 nm (947.7 nm) and an additional increase in the FWHM of 5 % (23 nm). An updated set of instrument-to-instrument comparisons can be seen in Fig. 10.

SAGE II is able to measure NO₂ by observing the difference in absorption between the two channels located at 452 nm (channel 5) and 448 nm (channel 6). Another look at exoatmospheric data (Fig. 6) shows that, like the water vapor channel, channel 6 may have had a change in spectral response prior to 1990 as a result of the aforementioned thermal vacuum testing incident. We chose a similar, but slightly different approach, to that used to correct the water vapor channel. Since NO₂ measurements are derived primarily from the differential extinction between the two channels, we chose to use one channel to calibrate the other, since the long-term variation in solar *I* zero showed channel 5 to be well behaved. The differential effective cross-sections between the two channels were compared to the differences in their exoatmospheric counts. It was assumed that any difference in the exoatmospheric counts of channel 6 relative to that seen in channel 5 over time were a result of a change in the band-pass of its filter. In this way, a time-dependent fit could be made that allowed the application of a differential cross-section correction to the effective cross-sections in

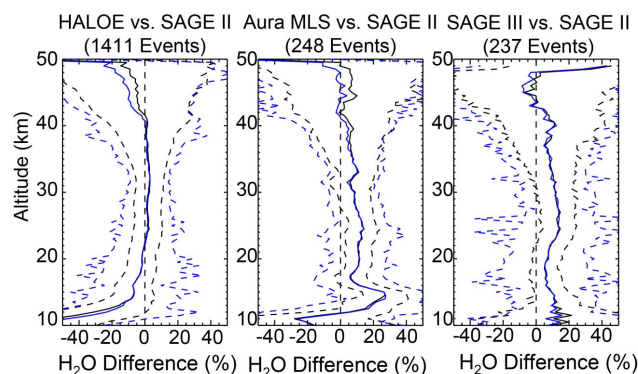


Fig. 9. A comparison of HALOE/Aura MLS/SAGE III minus SAGE II v6.2 water vapor products over their coincident events ($< 2^\circ$ lat, $< 10^\circ$ lon, < 2.5 h). For each plot, mean data are shown in black while median data are shown in blue (uncertainties are shown with dashed lines). Since HALOE was used to infer channel drift in v6.2, they are obviously in good agreement. However, there is an offset of about 10 % with both Aura MLS and SAGE III.

channel 5 to obtain the effective cross-sections in channel 6 for any time during the mission. The time-dependent model for the channel 6 spectral band-pass required the determination of the initial properties and the rate of change relative to the observed long-term anomalous *I* zero drift. To compute these, the SAGE II NO₂ stratospheric column abundances were compared to twilight NO₂ measurements at Lauder, New Zealand (Johnston and McKenzie, 1984). This procedure was applied in version 6.2, though only at one temperature (240 K). The dependence of SAGE II NO₂ retrieval on O₃ necessitated a repeat of this procedure after both the O₃ and NO₂ cross-section databases were changed. The difference in version 7.0 is the implementation of a temperature-dependent differential cross-section correction and the use of SAGE III NO₂ data to determine the starting spectral location of channel 6. The overall quality of SAGE II NO₂ measurements relative to the SAGE III NO₂ measurements remains mostly unchanged between versions (Fig. 11).

3 Transmission

The transmission algorithm combines measured limb-darkening curves with timing and pointing data to produce slant-path optical depths for each tangent altitude between 0.5 km and 100 km in 0.5 km increments. In order to do this, a number of physical and instrumental sources of variation must be compensated for.

The series of calculations and corrections for version 6.2 were done in a mostly linear fashion with some iterative calculation of transmission. It is now recognized that many of these corrections are dependent upon transmission and are thus updated in version 7.0 throughout these iterations. The algorithm begins by computing edges and pointing

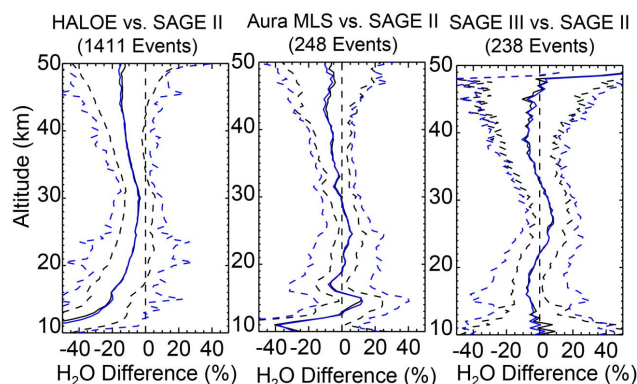


Fig. 10. A comparison of HALOE/Aura MLS/SAGE III minus SAGE II v7.0 water vapor products over their coincident events ($<2^\circ$ lat, $<10^\circ$ lon, <2.5 h). For each plot, mean data are shown in black while median data are shown in blue (uncertainties are shown with dashed lines). This shows the SAGE II v7.0 water vapor product to be in much better agreement with Aura MLS and SAGE III data sets than v6.2.

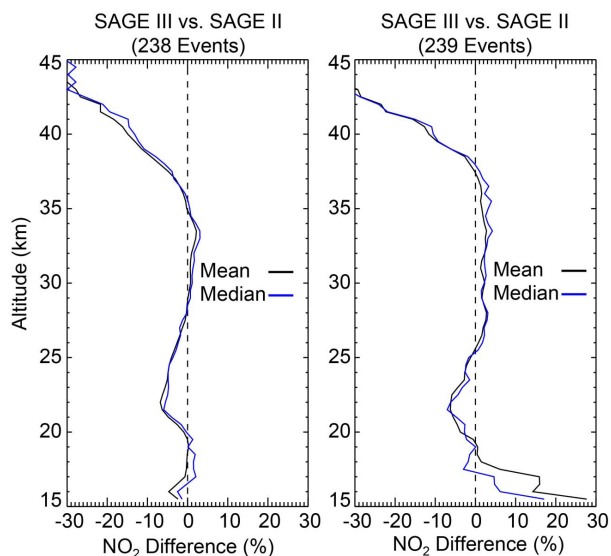


Fig. 11. SAGE III NO_2 minus SAGE II v6.2/v7.0 (left/right) over coincident events ($<2^\circ$ lat, $<10^\circ$ lon, <2.5 h). The overall quality of SAGE II NO_2 remains mostly the same between versions. The bias above 38 km is a source of ongoing study.

information (Sect. 3.1) to place each data packet on the face of the Sun, followed by computing a preliminary I zero (Sect. 3.2). From this, a preliminary transmission is created (avoiding sunspots) and the mirror calibration is computed (avoiding PMCs and accounting for Rayleigh scattering) (Sect. 3.4). The I zero curves are then updated (including a correction for the electronic transient (Sect. 3.3) if the event is a sunset). Once I zero curves have been obtained, a refined estimate of transmission is computed. The edge-time refinement algorithm (Sect. 3.5) is iterated for convergence to improve the point registration on the face of the Sun and




Fig. 12. Solar irradiance data measured by SAGE II for a typical sunset event showing both exoatmospheric and in-atmosphere scans. The shortening of scan duration as a function of flattening of the Sun due to refraction can clearly be seen. The alternating duration of each scan is a byproduct of spacecraft and orbital motions. Typically for sunset events, scans that move from the bottom to the top of the Sun (“up” scans) take less time than scans that move from the top to the bottom of the Sun (“down” scans).

a resultant transmission profile is computed (Sect. 3.6). If the event is a sunrise, the electronic transient correction is computed and applied. Lastly, a time-dependent I zero correction is performed and the entire process is reiterated so that the algorithm can take advantage of every correction available. The algorithm then proceeds to resample the transmission data to the standard grid and compute uncertainties (Sect. 3.6).

3.1 Edges and pointing

The first step in transmission processing requires that the limb-darkening curves (Fig. 12) be combined with pointing data to place every data packet accurately on the face of the Sun. First the limb-darkening curves in the 1020 nm channel are used to determine where the physical top and bottom edges of the Sun are by looking for the inflection points and the times associated with them. The timing of science packet data and ephemeris pointing data can then be accurately mapped to each other. This mapping is done by assuming that the rate of motion of the scan-mirror is constant during a scan so that each packet of data within a scan can be interpolated to a location on the face of the Sun.

This assumption becomes invalid when the bottom of the Sun is obscured by cloud or is below the limb of the Earth, as the calculated inflection point of the limb-darkening curve no longer correlates to the physical edge of the Sun and the calculated scan rate becomes biased high. The apparent rate of motion of the scan-mirror is a combination of the motion of

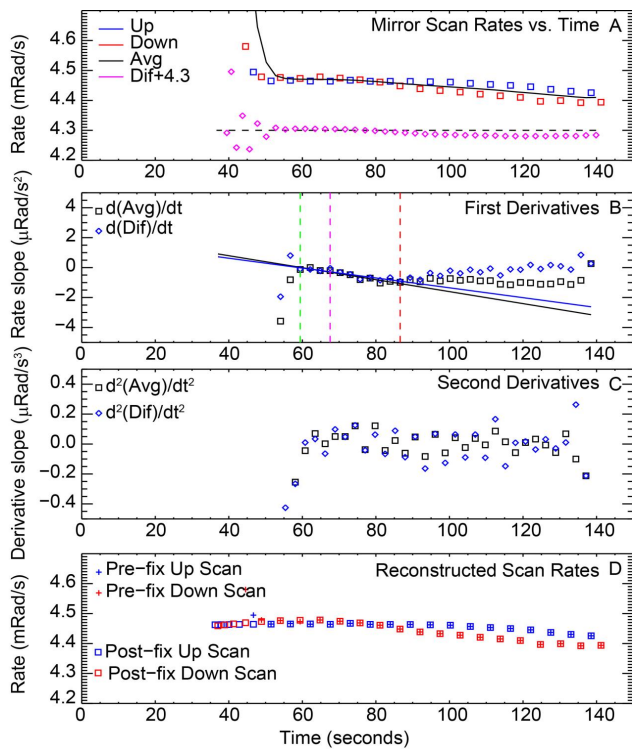


Fig. 13. Scan-mirror rate data for a sunrise event. Fits to the first derivatives of the average scan rate and difference between up and down scan rates are used to reconstruct scan rate data at low altitudes where scan rate data have gone bad. Only the scans just above where rate data go bad are used (between the green and red dashed lines in (B)). To avoid any programmatic errors with using the second derivative data to determine where rates have gone bad, an altitude upper limit of 60 km is used (dashed purple line). The comparison between original and reconstructed scan rates is shown in (D).

the scan-mirror itself and the orbital and attitudinal motions of the spacecraft. A look at the average scan rates and the difference between up and down scan rates for a typical event shows that they are generally well behaved and slowly varying with time (Fig. 13a). The derivatives of these quantities show linearity with respect to time (Fig. 13b). However, at any time during the event, the attitude actuators on the spacecraft (these keep the spacecraft in a desired orientation) can turn on or off and cause an abrupt shift in these derivatives. While attitude control maneuvers do not affect retrieved scan rates due to the fact that they affect the retrieved top and bottom edge times equally, they do account for the kink in Fig. 13b. The linearity of derivatives with respect to time allows for a fit of these quantities to data just above where the rate data begins to go bad, which is characterized by large values in the first or second derivatives (Fig. 13c). These fits are then used to reconstruct low altitude rate data (Fig. 13d) in order to compensate for scans with bad edge calculations resulting from the bottom of the Sun being occluded by cloud or the limb of the Earth.

Once the algorithm has good rate data for each scan, the proper tangent altitude and position on the face of the Sun are calculated for each packet. Given the problem of combining refraction effects derived from a spherical model dependent upon distance from the center of the Earth and an oblate Earth model dependent upon distance from the surface of the Earth, the algorithm uses an iterative scheme optimized for rapid convergence to refine the tangent point altitudes. During the development of both the SAGE III version 4.0 and SAGE version 7.0 algorithms, the algorithmic uncertainty in determining the tangent point altitude was determined to be better than 20 meters. Lower in the atmosphere (in the troposphere), where refraction effects can become large, the uncertainty in the tangent point altitude is dominated by uncertainties in the meteorological data. Each packet of data, in each channel, is subsequently assigned a time, tangent altitude, position on the face of the Sun, scan-mirror elevation position, and photodiode count value.

3.2 *I* zero

In order to create transmission profiles, the algorithm calculates ratio measurements made by the instrument looking at a particular point on the Sun through the atmosphere to the same point seen above the atmosphere. Thus one of the first things the transmission algorithm does is create a standard exoatmospheric limb-darkening curve (*I* zero curve) ratio measurement to map other scans to. The instrument generally collects between 10 and 20 exoatmospheric scans, so these are combined into a pair of *I* zero curves, one for up-scans and one for down-scans, which are then interpolated to a fine grid in position on the face of the Sun (~ 1000 points in version 7.0 versus ~ 100 points in version 6.2). As the *I* zero curve is mapped onto each scan (including the exoatmospheric or *I* zero scans), an edge time refinement (see Sect. 3.5) is made for the *I* zero scans. A new addition in version 7.0 is the introduction of a time-dependent *I* zero correction, which was first introduced in SAGE III version 4.0. A time-dependent *I* zero correction benefits high altitude scans by helping to correct for apparent rotation of the scan track across the face of the Sun due to orbital motions (Burton et al., 2010).

Occasionally the instrument would fail to acquire more than one or two exoatmospheric scans. This happened sporadically through the lifetime of the instrument, but most notably during the so called “short event” period (from mid-1993 to mid-1994), when a battery problem on the spacecraft reduced the operational scan time of the instrument, causing sunset events to begin later than normal and sunrise events to end earlier than normal. Many of the corrections applied to the data (e.g., mirror calibration and “thermal shock”) require a minimum number of *I* zero scans (i.e., data over a large range of exoatmospheric altitudes). In prior versions, the absence of this data would often result in anomalous events that were not screened (i.e., dropped) by the algorithm and

required the user to manually screen them out (Wang et al., 2002). In version 7.0, all events without a minimum number of I zero scans are dropped from processing and flagged accordingly.

3.3 Other corrections

Once an I zero curve is established (and mapped to each scan), a preliminary transmission can be computed for each scan. This then allows the calculation of a few corrections that need to be made. An initial correction for Rayleigh attenuation is done (for the benefit of other corrections) and a polar mesospheric cloud (PMC) detection routine is run (Burton and Thomason, 2000). While no overall correction is made in the presence of PMCs, the use of data within PMCs will be avoided when determining the mirror calibration (see Sect. 3.4). A sunspot detection routine is also run and measurements inside of sunspots are flagged so that they can be filtered later. It was found that the sunspot detection routine in version 6.2 would often overcompensate in its sunspot filtering (i.e., it would often omit all data from the start of a sunspot to the edge of the Sun) and a more robust algorithm has been introduced for version 7.0, which identifies sunspots by their characteristic shapes in both the first and second derivatives of the limb-darkening curves. In order to avoid false positives resulting from noise in the limb-darkening curves, the derivatives are normalized and the magnitude and shape of both derivatives are used simultaneously to identify sunspots.

As the SAGE II instrument began taking data during each event, a small transient was observed in channels 5 and 6. This so-called thermal shock needed to be accounted for in the data processing. For sunset events, this is relatively easy as the data are fairly homogenous (exoatmospheric scans) and any time dependency is easily recognized and removed. For sunrise events, however, the transient occurs while the instrument is looking through the attenuated atmosphere and the correction is necessarily at low altitudes. In this case, scan-to-scan variations in the channel 5 and 6 differential extinction at a given altitude, which are correlated in time across multiple scans, are examined. By looking at the ratio of channel 6 to channel 5, which is used for the NO_2 retrieval, the algorithm determines the rate of change of differential extinction, which is then integrated to produce a correction that is applied to these two channels (Fig. 14). The impact on the transmission in these channels is on the order of 0.25 %, whereas the impact on other channels (were it to be applied) would be no more than 0.05 % (SPARC/IOC/GAW, 1998).

3.4 Mirror calibration

SAGE operates by using a scan-mirror to move the instrument field-of-view up and down across the solar disk (normal to the Earth's limb). The reflectivity of the scan-mirror varies slightly with the angle of incidence and requires calibration.

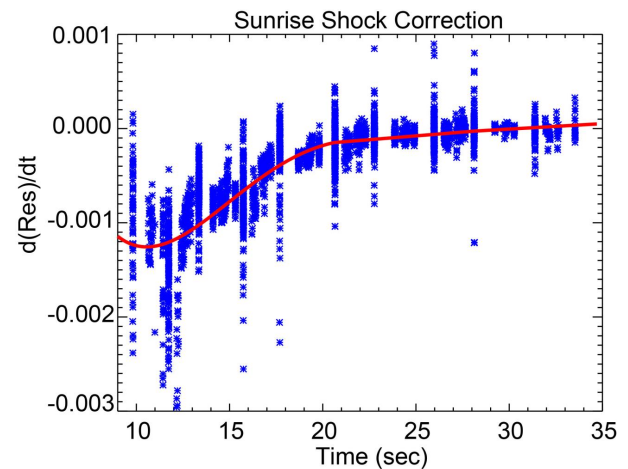


Fig. 14. The sunrise shock correction fits the differential extinction between channels 5 and 6 with an exponential and computes the residuals. The residuals within an altitude bin are then analyzed, computing the rate of change of the residuals as a function of time. The blue asterisks in the plot denote the computed derivatives while the red line denotes the fit. Once computed, this fit is integrated to derive a correction factor to apply to the originally computed differential extinctions.

This is not an absolute reflectance calibration as only the relative change in reflectivity with angle is required. In version 6.2, this calibration was determined through a quadratic fit to exoatmospheric transmission data as a function of altitude and was performed only once at the end of transmission processing. In version 7.0, this quadratic fit is to exoatmospheric transmission data as a function of the elevation angle of the scan-mirror and is performed in the overall iterations of transmission processing. In each case, the mirror calibration is applied as a multiplicative correction to remove curvature in the high altitude transmission data, which should theoretically be a constant value of 1 with some instrument noise. Since smaller angles (higher altitudes) are used for the fit, an extrapolation of the correction term to larger angles (lower altitudes) is necessary to correct attenuated data. An example of the mirror correction can be seen in Fig. 15, demonstrating that the angular dependence of the reflectivity of the scan-mirror is on the order of 0.5 %.

3.5 Edge-time refinement

The edge-time refinement algorithm is designed to minimize any biases in the limb-darkening curves created by slight errors of the initial calculation of the edge times of each scan. The measured solar intensity of a scan (I) in a given channel (λ) is a function of both position on the face of the Sun (p) and altitude (z), which are themselves both functions of time (t), and can be written as

$$I(\lambda, p(t), z(t)) = I_0(\lambda, p(t)) T(\lambda, z(t)) + \varepsilon(\lambda, p(t), z(t)), \quad (1)$$




Fig. 15. SAGE II v7.0 mirror correction for the 452 nm channel on the very first iteration of a sunset event. The correction is a quadratic fit to raw high-altitude transmission as a function of scan-mirror elevation angle. The black dots show the raw transmission values and the blue line shows the correction factor. The corrected transmission (red dots) are raw transmission divided by the correction factor. The plot on the left shows transmission data points and correction factors versus scan-mirror elevation angle while the plot on the right shows the same points plotted against the corresponding altitude. The correction factors are extrapolated outside of the range of the fitted transmission data points for use with low altitude transmission data.

where I_0 is the I zero limb-darkening curve, T is the slant-path transmission, and ε is the error in the measurements or estimates. Since there is some inherent uncertainty in the calculated edge times, Eq. (1) can realistically be rewritten as

$$I(\lambda, p(t + \varepsilon_t), z(t + \varepsilon_t)) = I_0(\lambda, p(t)) T(\lambda, z(t)) + \varepsilon(\lambda, p(t), z(t)), \quad (2)$$

where ε_t is related to uncertainties in the edge times. There are two cases that can be considered separately. The first case is for high altitude measurements where $T(z) = 1$, and Eq. (2) can be expanded into

$$\begin{aligned} I(\lambda, p(t + \varepsilon_t)) &= I(\lambda, p + \varepsilon_p) \\ &= I_0(\lambda, p) + \frac{dI_0(\lambda, p)}{dp} \varepsilon_p \\ &= I_0(\lambda, p) + \frac{dI_0(\lambda, p)}{dp} (c_1 + c_2 p), \end{aligned} \quad (3)$$

where I_0 is the estimate for the I zero curve and c_1 and c_2 are the linear shift and stretch coefficients, respectively. The correction is a linear function because position on the face of the Sun is linearly mapped to time. The derivative term is calculated using finite differences and a multiple linear regression technique is used to obtain the shift and stretch terms, which yield the correction to the edge times.

The second case involves measurements in the atmosphere where $T(z)$ can no longer be ignored. In this case, Eq. (2) can

be expanded into

$$I(\lambda, p(t + \varepsilon_t), z(t + \varepsilon_t)) = I_0(\lambda, p(t)) T(\lambda, z(t)) + \frac{d}{dt} (I_0(\lambda, p(t)) T(\lambda, z(t))) \varepsilon_t, \quad (4)$$

$$I(\lambda, p(t + \varepsilon_t), z(t + \varepsilon_t)) = I_0 T + \left(I_0 \frac{dT}{dz} \frac{dz}{dt} + T \frac{dI_0}{dp} \frac{dp}{dt} \right) (c_1 + c_2 t), \quad (5)$$

where c_1 and c_2 are again the linear shift and stretch (S&S) coefficients. Following the same procedure, the correction to the edge times is obtained. Since the pair of edge times apply equally to all spectral channels, the derivatives are performed at each wavelength and used together in the regression calculation. The shift and stretch algorithm has the effect of removing spectrally and vertically correlated noise in the transmission profile, typically at low altitudes (Fig. 16).

3.6 Normalization, gridding, and uncertainties

In order to calculate the various corrections already outlined, the algorithm iteratively computes transmission values for each packet. In order to mitigate effects remaining from residual edge time uncertainties, the outer 10 % of the solar disk is omitted. Omitting the edges of the Sun is also meant to compensate for another correction that is currently not applied, namely a correction for the expansion of the instrument's field-of-view on the face of the Sun as a function of refraction. The impacts of this effect are most notable at the edges of the Sun, where curvature in the limb-darkening curve is largest. After the algorithm has undergone several iterations of calculating transmission and applying corrections, it computes the final transmission profile. A running median filter is applied through altitude sorted transmission packets to minimize the impact of strong outliers, and then smoothed with a boxcar average. The filtering process is performed in altitude and the filtering/smoothing parameters correspond to 1.0 km so that all transmission data meet the Nyquist sampling criteria for a 0.5 km gridded profile. An example of an intermediate transmission profile and a final transmission profile is shown in Fig. 17. The transmission data are then interpolated to the 0.5 km grid, and the variance of the fit with respect to the raw data in each bin is used to compute the uncertainty estimate. In version 6.2, the statistics of this fit were used for the uncertainty in the transmission value in each altitude bin. In version 7.0, we have incorporated an additional uncertainty term, namely a calculated uncertainty in the original I zero curve, meant to account for variations between the exoatmospheric scans used to create the I zero curve and the resulting I zero curve itself. These minor variations between each exoatmospheric scan are highly correlated with position on the solar disk, but have no discernible time dependency (i.e., they are not detected and filtered by the time-dependent I zero correction). It is believed that these variations represent physical features




Fig. 16. The effect of the shift and stretch (S&S) algorithm on 1020 nm transmission during the first iteration for two separate SAGE II events. The black marks/lines denote pre-S&S transmission (shifted to the left by 0.2 in transmission for ease of viewing) while the blue marks/lines denote post-S&S transmission. The shift and stretch algorithm helps to remove a lot of the noise associated with inaccurate edge times.

in the Sun's photosphere (e.g., granulation) combined with the apparent rotation of the solar disk during the event, as opposed to instrumental noise. These variations manifest themselves as low amplitude oscillatory patterns in high altitude transmission (on the order of 0.1 %) that are periodic in altitude (due to their high correlation with surface features on the Sun) and correlated between channels (Fig. 18). An extension of the time-dependent I zero algorithm to compensate for this effect is in development. Lastly, all slant-path transmission profiles are converted to slant-path optical depth profiles.

4 Vertical profiles of individual species

The inversion algorithm takes slant-path optical depth profiles and, along with other data such as P/T data, separates them into species-specific slant-path optical depth profiles before finally inverting them into vertical profiles of O_3 and NO_2 number densities, H_2O volume mixing ratio, and aerosol extinction. While much of this follows the same basic procedure described in Chu et al. (1989), there are some important differences between versions 6.2 and 7.0, some of which include more subtle aspects of the algorithm. As such, this section will generally review the entire inversion process. The uncertainties throughout this process are propagated from uncertainties in the neutral density and derived transmission profiles.




Fig. 17. The normalization and fitting process for most iterations (left) and the final iteration (right) of the algorithm. Black marks denote transmission values for each packet while the blue lines denote the fitting process.

4.1 Basic procedures

The viewing geometry of SAGE is such that at each tangent height the instrument looks through the atmosphere, it is also looking through a slant-path column of air that incorporates all of the tangent heights above it. Accounting for refraction, a triangular path-length matrix is computed. The slant-path total column at each tangent height, derived from a matrix multiplication of the path-length matrix and density profile, is thus comprised of the sum of partial slant-path columns from all overlying 0.5 km thick layers. This simple matrix operation allows for an “onion peeling” process to be performed later for inverting a species' slant-path optical depth profile to the species density profile.

The first step in the retrieval of vertical profiles of individual species is to account for and remove the contributions of molecular (Rayleigh) scattering in all channels and O_2 - O_2 absorption in a subset of channels. O_2 - O_2 cross-sections ($\text{cm}^5 \text{ molecule}^{-1}$) are scaled with density and both O_2 - O_2 and Rayleigh effective cross-sections ($\text{cm}^5 \text{ molecule}^{-2}$) are combined with the slant-path total column to convert to slant-path optical depths, which are then subtracted from each channel. The largest source of uncertainty in the retrieved profiles of individual species comes from uncertainty in the contribution from molecular scattering, which itself originates from uncertainty in the temperature profile.

4.2 Species separation

The species separation in SAGE II is performed using five of the seven channels simultaneously (1020, 600, 525, 452, and 448 nm) and is separated into three altitude regions: altitudes where some of the five channels are no longer available,

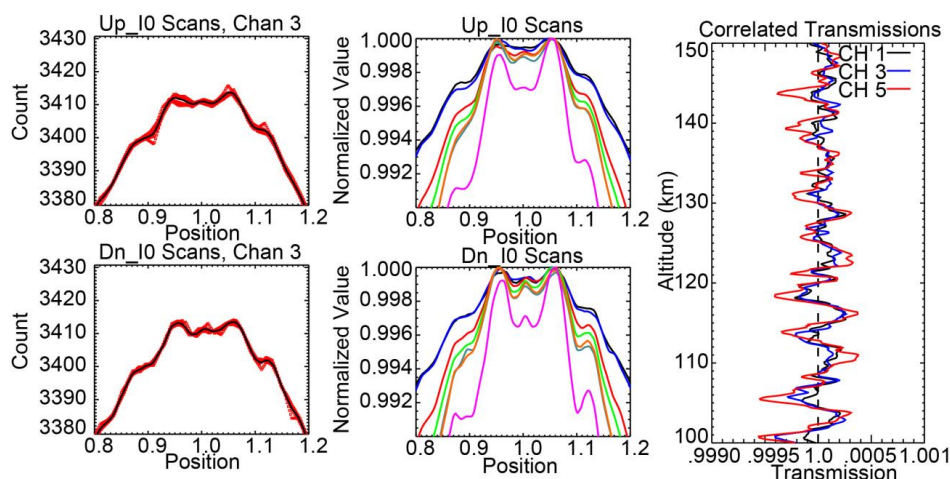


Fig. 18. The plots on the left show the creation of an I zero curve for both up (top) and down (bottom) scans. Each I zero curve (black line) is created from multiple exoatmospheric scans (red marks). These plots are zoomed in on the center 20 % of the Sun scan to show how there are both minor variations in intensity across the face of the Sun as well as some inherent uncertainty in the resulting value of the I zero curve. The plots in the middle show normalized up (top) and down (bottom) I zero curves for each channel to show how these spatial variations are strongly correlated across channels (channels 1–7 are discernible from outermost to innermost). It is very likely that this variation is representative of physical variation on the Sun (e.g., solar granulation). The plot on the right shows high altitude transmission for several channels. The repetitive oscillatory patterns that are strongly correlated between channels are the direct result of this unfiltered physical variation. For now, this is not compensated for in the final transmission profiles but, as a new feature in version 7.0, is accounted for in the uncertainty estimates.

altitudes in which all five channels are available, and altitudes above which it is believed there is no aerosol extinction. These five channels are used because the poor quality of the 386 nm channel prevents its use in the broad retrieval. Significant absorption by water vapor occurs only in the 940 nm channel and is inverted in a separate process.

The first step is to find the highest extent to which one of the five channels is not available (i.e., no valid or non-fill data) and begin just above that. The residual slant-path optical depth (hereafter abbreviated as OD) in each channel consists of contributions from O_3 , NO_2 , and aerosol, though at longer wavelengths (particularly 1020 nm) the contribution from gas species absorption becomes vanishingly small. At this point we have measurements with seven unknowns: O_3 , NO_2 , and aerosol at each wavelength. We solve this set of equations using a least-squares solution where we approximate the aerosol contribution at 600 and 448 nm as the linear combination of aerosol at 1020, 525, and 452 nm. The full set of species separation equations to be solved can be expressed as

$$c_1(\lambda_i) OD_{Aer}(\lambda_1) + \frac{\sigma_{O_3}(\lambda_i)}{\sigma_{O_3}(\lambda_3)} OD_{O_3}(\lambda_3) + c_2(\lambda_i) OD_{Aer}(\lambda_4) \\ + c_3(\lambda_i) OD_{Aer}(\lambda_5) + \frac{\sigma_{NO_2}(\lambda_i)}{\sigma_{NO_2}(\lambda_5)} OD_{NO_2}(\lambda_5) = OD(\lambda_i)$$

where $i = \{1, 3, 4, 5, 6\}$ is the channel number, $\sigma(\lambda_i)$ is the effective cross-section of the stated species at the given channel, $OD(\lambda_i)$ is the slant-path optical depth of the stated channel, and the channel specific aerosol, O_3 , and NO_2 ODs are

the unknowns. Most of the coefficients for this process (c_1 , c_2 , and c_3) are simply “1s” and “0s” (e.g., for channels 1, 4, and 5), and those that are not are determined using an ensemble of single mode log-normal size distributions of sulfate aerosol at stratospheric temperatures, though, in practice, composition is of secondary importance. The ensemble of log-normal size distributions spans the observed wavelength-dependence of the aerosol spectra and is used to estimate the relative dependence of extinction in channels 3 and 6 as a function of the values in channels 1, 4, and 5. The ensemble consists of a large range of mode radii and widths that effectively span the observed relationships of the dependent and independent aerosol channels. Ultimately this is simply a means of interpolating between nominal aerosol channels as, given the log-normal ensemble extinction estimates at 1020, 525, and 452 nm, the range of values possible at 600 or 448 nm is small. Using simple linear processes, the coefficients that relate the dependent extinction values to the three independent extinction values are robust, and residuals of the fit provide at least some information regarding uncertainty in clearing aerosol from each channel.

Once O_3 and NO_2 ODs are determined, their relative contributions in the 940 and 386 nm channels can be removed and aerosol OD in the 386 nm channel is retrieved as a residual. Aerosol OD in the water vapor channel is calculated from the 525 and 1020 nm aerosol values, using separate weighting coefficients determined from the use of the 525 to 1020 aerosol OD ratio. Once the aerosol contribution in the water vapor channel is determined, the actual water vapor

OD is calculated as a residual. This process separates the various species from lower altitudes (typically in the middle to upper troposphere) up to some maximum altitude (typically set to 75 km).

While this retrieval works very well, it typically suffers from attempting to retrieve species at higher altitudes where extinction values are near detection limits and noise becomes the dominant signal. This noise is a carryover of the channel-correlated solar structure noise first discussed in Sect. 3.6, which does not dampen at higher altitudes. Much of this noise ends up being interpreted by the algorithm as aerosol and detrimentally affects the simultaneous retrieval of all species. To compensate for this in version 6.2, a separate retrieval scheme was used that uses only 4 channels (600, 525, 452, and 448 nm) and assumes that no aerosol is present to calculate O_3 OD in the 600 nm channel and NO_2 OD in the 448 nm channel. Water vapor was again treated as a residual in the 940 nm channel. In lieu of more adaptive methods, this process began at 40 km up to some maximum altitude (typically 75 km). Thereafter, the 5-channel retrieval was transitioned into the “no aerosol” retrieval between 40 and 45 km. The version 7.0 algorithm utilizes the data to determine how to transition into regions where the inclusion of aerosol in the retrieval is no longer necessary, the methods of which are outlined later in this section.

For altitudes below the 5-channel retrieval, there is no longer valid data in the 448 nm channel and thus NO_2 cannot be simultaneously retrieved. Instead, the NO_2 OD profile from the 5-channel retrieval is inverted to get extinction values and the algorithm reconstructs ODs at lower altitudes by assuming the NO_2 mixing ratio is zero. The OD contribution from NO_2 at lower altitudes is then removed from all channels. With NO_2 removed, the algorithm begins working from the bottom of the 5-channel retrieval and moves down. It first uses a 4-channel retrieval (1020, 600, 525, and 452 nm) so long as there is valid data and then transitions to a 3-channel retrieval (1020, 600, and 525 nm) when necessary. The main retrieval algorithm stops if there is no longer any valid data in any of these three channels. Aerosol extinction in the 1020 nm channel is thereafter retrieved until no further data exists. These data are then used to estimate the contribution of aerosol in the 940 nm channel so that water vapor can again be calculated as a residual.

This entire process goes through two iterations. The first iteration uses a default set of weighting coefficients to determine the contribution of aerosol at 600 and 448 nm. In the second iteration, we use the measured 525 to 1020 nm aerosol extinction ratio to select sets of coefficients determined using the fits to the ensemble of aerosol spectra for values around the observed value. This accounts for a small degree of nonlinearity observed in the fits to 600 and 448 nm aerosol extinction. In reality, this is a distinctly second order correction but seems to reduce the sensitivity of the quality of the ozone data product to aerosol, particularly when aerosol levels are high. However, for the same reasons that version

6.2 has a “no aerosol” retrieval, at altitudes near and above 40 km, the 525 to 1020 aerosol OD ratio can begin to vary wildly through non-physical numbers. This has a detrimental effect on the second iteration retrievals, producing extremely “noisy” data above 35 km. In version 7.0, once the retrieved aerosol OD in each channel drops below a predetermined threshold, a nonlinear least squares fit is made to the data in the form of an exponential decay curve. Above the altitude where these fits drop below the amplitude of the noise, the fits are used for the 525 to 1020 ratio instead of the actual data. The second iteration retrievals then have far more realistic weighting coefficients to work from.

While the use of fits to determine the 525 to 1020 ratio does improve the quality of the 5-channel retrievals, particularly above 35 km, it can still suffer from the same limitations that necessitated the use of a “no aerosol” retrieval above 40 km, namely the fact that the aerosol ODs at higher altitudes can still become non-physical and detrimentally affect the simultaneous retrieval of all species. Typically this is a byproduct of retrieved non-physical aerosol OD values in the shorter wavelengths at higher altitudes detrimentally impacting the retrieved values in longer wavelengths. This had the tendency (in version 6.2) to result in O_3 and NO_2 values that were biased high at altitudes above ~ 40 km. To compensate for this, in version 7.00, the algorithm looks only at the 600 nm channel to retrieve ozone in these altitude regimes. When the fit to the 600 nm aerosol OD drops below a certain threshold, the algorithm subtracts out the fit (which is generally below the “noise level”) and inverts only the remaining 600 nm OD to retrieve ozone. A more adaptive algorithm is being developed for retrieving NO_2 at higher altitudes though, for now, the high bias persists from version 6.2.

The water vapor retrieval has benefited greatly from the use of fits to the aerosol optical depths to determine the 525 to 1020 ratio. In version 6.2, once either the 525 or 1020 aerosol contribution became negative, the algorithm assumed there was no longer any aerosol contribution above that altitude in the water vapor channel (mainly because the weighting coefficients could not be determined from a negative ratio). However, it is possible, due to noise, for the aerosol contribution to become negative but then become some non-negligible positive number again. This should have manifested itself as noise in the data, but was instead being artificially removed. By using the fits to the aerosols to determine the 525 to 1020 ratio, weighting coefficients could be applied to the real data (be it positive or negative) through the entire retrieval range. This has benefited the retrieval greatly as, after the removal of O_3 and NO_2 , the aerosol contribution to the water vapor channel can still be a significant fraction of the remaining signal in some cases.




Fig. 19. Coincident event mean water vapor mixing ratio profiles between SAGE III (blue) and SAGE II (black) (v6.2 left and v7.0 right). In v6.2, SAGE II water vapor mixing ratio profiles approached asymptote towards large values near the top of the retrieval, a trait that SAGE III v4.0 also exhibits. In v7.0, this has been corrected. The overall change in bias between v6.2 and v7.0 comes from a change in ozone spectroscopy and a modification of the water vapor channel filter response towards results in better agreement with SAGE III and Aura MLS.

4.3 Inversion

After species separation, water vapor is the first retrieved species. The slant-path optical depth data are first converted to extinction and then smoothed to help mitigate noise in the weak signal. The algorithm to retrieve water vapor mixing ratio remains mostly unchanged from Chu et al. (1993) with one small exception. Originally the process began at 50 km and worked down. The algorithm requires a small range of altitude at the top of the retrieval process to establish the abundance and scale height of the H_2O profile as a boundary condition. Improvements in the version 7.0 transmission allowed the starting altitude to be moved up to 60 km. This, combined with the aforementioned aerosol fit method, has greatly improved the water vapor product above 40 km. In version 6.2, water vapor mixing ratio profiles would suffer from a characteristic “hook” towards unrealistically large values nearing 50 km. This can also be seen in SAGE III version 4.0 water vapor mixing ratio profiles. With these modifications to the retrieval, this “hook” has been removed (Fig. 19).

Prior to version 7.0, once water vapor was retrieved at the end of the first iteration, its contribution to the 600 nm channel was removed and the second iteration began. This had an inconsequential impact on ozone above the hygropause. However, below the hygropause, SAGE II ozone retrievals have consistently been biased low when compared with other data (Wang et al., 2002). This feature was turned off in

SAGE III version 4.0, due to uncertainty of the quality of the relative spectroscopy used between the water vapor channel and the 600 nm channel, with beneficial results (Wang et al., 2006). For the same reason, this feature has also been turned off in SAGE II version 7.0. It is important to note, however, that this has not completely corrected the low bias in SAGE II tropospheric ozone data.

Once the algorithm has run through both species separation iterations, it has produced vertical water vapor volume mixing ratio and NO_2 number density profiles. All that remains is to invert the O_3 and various aerosol slant-path column optical depth profiles into vertical extinction profiles. As with NO_2 , the O_3 vertical extinction profile is then converted to a vertical number density profile by simply dividing by the effective cross-section. The inversion technique used in version 7.0 is different from that used in version 6.2, which utilized Twomey’s modification of Chahine’s algorithm (Chahine, 1972; Twomey, 1975) to retrieve extinction values and a simple onion-peeling technique to retrieve uncertainty estimates. The Twomey–Chahine algorithm was found to have several undesirable behaviors. It would not allow negative values and therefore introduced a positive bias in regions of the density (extinction) profile where the signal-to-noise ratio was small, typically at the higher altitude end of the retrieved profile. It also systematically approached the solution from one direction and stopped once the tolerance criterion was met, producing another form of bias as a result. Lastly, its use introduced discontinuities in the profile when the slant-path extinction fell below a preset value and vertical smoothing was activated. Given the high quality of the SAGE II version 7.0 transmission profiles and the algorithmic limitations of the Twomey–Chahine inversion method, it was replaced entirely with onion-peeling in version 7.0.

Lastly, with all primary data products computed, the aerosol extinction in the 525 and 1020 nm channels are used to compute some physical parameters to characterize aerosol at each altitude. The methods outlined in Thomason et al. (2008) are used to determine the effective radius and surface area density of aerosol particles.

5 Results of version 7.0

This paper describes the SAGE II version 6.2 and version 7.0 algorithms. Prior versions of SAGE II data products have been well validated (e.g., Wang et al., 2002) and included in numerous international assessments (e.g., WMO, 2011). Several of the version 7.0 changes affect the quality of these data products and a brief assessment of the differences seen in the version 7.0 data products follows.

5.1 Event comparisons

Figure 20 shows several comparison plots of ozone. The changes in ozone from version 6.2 to version 7.0 come

almost entirely from the change in spectroscopy, resulting in a nearly uniform decrease in concentration of $\sim 1.5\%$ in the stratosphere. The magnitude of the offset increases to about 2% above 40 km due to the removal of the “no aerosol” method of retrieval. The large increase in ozone below 10 km is a result of the removal of the water vapor correction to ozone. The difference between the mean and median below 20 km is a result of using onion-peeling for inversion instead of Twomey–Chahine, as the outliers in version 6.2 were biased positive whereas in version 7.0 they can take on negative values, though overall the same number of outliers exist. While the changes move SAGE II ozone values further from SAGE III concentrations by about 1.5% despite using the same spectroscopy, the changes in the retrieval method make the offset vary less with altitude.

Given the diurnal nature of NO_2 , sunrise and sunset events are compared separately. Figure 21, left and middle panels, show comparison plots for sunset and sunrise NO_2 , respectively. Due to changes in spectroscopy, the overall concentration of sunset NO_2 has dropped by 5–10% in the mid-stratosphere. The large decrease below 20 km is again a result of using onion-peeling instead of Twomey–Chahine for inversion. As previously mentioned, sunset NO_2 values are biased high above 40 km and an algorithm to better retrieve NO_2 at these higher altitudes is in development. Net changes in sunrise NO_2 are an amalgamation of changes made in transmission, spectroscopy, species separation, and inversion. Sunrise NO_2 remains somewhat of a research product as quantifying the impact on sunrise NO_2 data quality is hampered by the fact that insufficient high quality sunrise NO_2 measurements exist for comparison during the time period where SAGE II measured sunrise NO_2 (all comparisons with SAGE III are sunset events due to power problems late in the mission, forcing operation at half duty cycle). Looking at the sunset/sunrise NO_2 ratio, however, reveals that the data are more consistent in the mid-stratosphere in version 7.0 than in version 6.2 (Fig. 21 right panel).

While no changes have been made to the retrieval of aerosol extinction in version 7.0 specifically, the data in various channels are impacted by the changes made to the spectroscopy and the technique used for inversion (Fig. 22). As a result of using onion-peeling instead of Twomey–Chahine for inversion, aerosol extinction tends to decrease more quickly at higher altitudes, as opposed to asymptoting to some positive non-zero value as in previous versions. The overall data quality through the mid-stratosphere has remained mostly unchanged for aerosol extinction in the 1020 and 385 nm channels. The change in spectroscopy has had a large impact on the comparisons between SAGE II and SAGE III aerosol extinction in the 525 and 452 nm channels. SAGE II version 7.0 aerosol extinction at 525 nm is in much better agreement with SAGE III version 4.0 as compared to SAGE II version 6.2, whereas the opposite is true for aerosol extinction at 452 nm. We have transitioned the aerosol derived products including surface area density

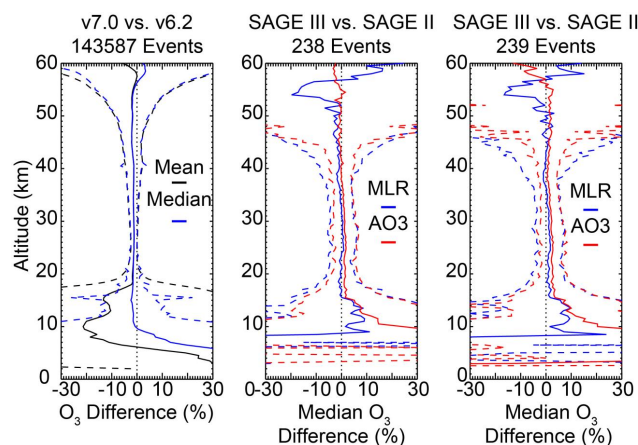


Fig. 20. Comparison plots of ozone. The panel on the left shows SAGE II v7.0 minus v6.2 for all events, excluding the time period of heavy aerosol interference from the Mount Pinatubo eruption. The other two panels show SAGE III v4.0 minus SAGE II v6.2 and v7.0 (middle and right, respectively) over their coincident events. The comparison of each of SAGE III’s main ozone data products is shown. MLR ozone (retrieved by a multiple linear regression technique) is represented in blue while aerosol ozone (retrieved in a fashion similar to SAGE II) is represented in red.

(SAD) and effective radius (R_{eff}) from the technique outlined in Thomason et al. (1997) to a more robust method developed in Thomason et al. (2008). Since there is some concern about the newer technique’s performance at low 525 to 1020 nm aerosol extinction coefficient levels, the new operational method transitions from the 2008 method for ratios above 2.0 to the old method for ratios below 1.5 with a linear mix in between. As a result, aerosol products do not change significantly during the post-Pinatubo period but change substantially during the clean period, particularly after 1998. The change in this period can be seen in Fig. 23 where the SAD has increased by 50% throughout the lower stratosphere and the R_{eff} has decreased by about 10%.

The process of retrieving water vapor and the impact on data quality has been discussed throughout this paper. As mentioned in Sect. 2.4, the water vapor channel filter spectral location was shifted a total of $+12.7$ nm and the FWHM was increased by 15% from the original location. This filter location was determined by comparing SAGE II and SAGE III water vapor data, albeit prior to the inclusion of the use of MERRA meteorological data in the retrieval. We have since revisited this problem and come to the conclusion that, in addition to the drift of the water vapor channel filter spectral response, the relative ozone spectroscopy used to remove ozone from the water vapor channel may be incorrect. The reasoning behind this is illustrated in several plots shown in Fig. 24. During the course of determining the best location to move the water vapor channel spectral response, it was noted that very good agreement could be reached in altitude regions of low ozone without adjusting the FWHM

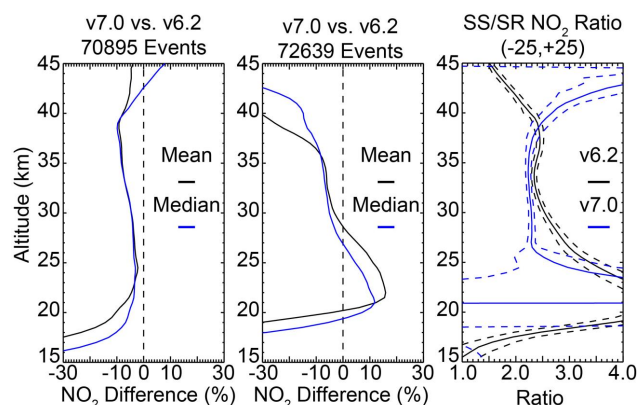


Fig. 21. Comparison plots of NO_2 . The plots at the left and middle show SAGE II v7.0 sunset and sunrise NO_2 , respectively, minus SAGE II v6.2 sunset and sunrise NO_2 , respectively, for all events, excluding the time period of heavy aerosol interference from the Mount Pinatubo eruption. The plot on the right shows the mean NO_2 sunset over sunrise ratio for all events (again excluding Pinatubo) in the tropics for v6.2 (black) and v7.0 (blue). Results are similar in other latitude bands.

of the channel. However, removing the remaining “ozone-like” signal from water vapor required nearly doubling the FWHM of the channel. It is well understood that some level of uncertainty exists in laboratory experiments to retrieve temperature-dependent ozone absorption cross-sections, particularly in the Wulf bands (Bogumil et al., 2003). Since the release of the SCIAMACHY V3 ozone cross-section database, several other ozone cross-section databases have been released that show significant changes in the Wulf bands relative to the Chappius (e.g., Chehade et al., 2013; Serdyuchenko et al., 2011). While we are hesitant to adopt a new cross-section database until it has been validated, SAGE II data suggests that the relative ozone spectroscopy in the SCIAMACHY V3 database could be off by on the order of 10 %. We have identified many possible combinations of changing both the filter location and the relative ozone spectroscopy in order to minimize differences with SAGE III water vapor. However, any change to the ozone spectroscopy creates a coupled problem when comparisons are made with SAGE III water vapor, as the same spectroscopy would have to be adopted by SAGE III as well. As it currently stands, the water vapor product in SAGE II version 7.0 is in much better agreement with SAGE III version 4.0 than was SAGE II version 6.2. However, we are still not satisfied with the result and this issue remains a topic of further study.

5.2 Time series analysis

Herein we present a new way of fitting SAGE II data for use with time series analysis; namely, we fit the entirety of the data at a single altitude simultaneously using the dates and latitudes of the measurements as they were made (i.e., no

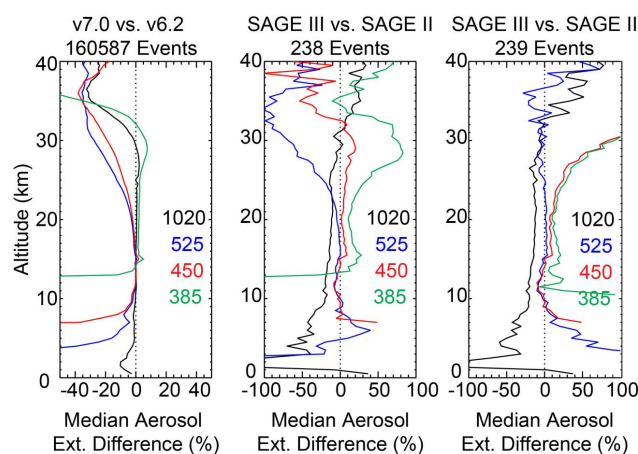


Fig. 22. Comparison plots of aerosol extinction. The panel on the left shows SAGE II v7.0 minus v6.2 for all events. The other two panels show SAGE III v4.0 minus SAGE II v6.2 and v7.0 (middle and right, respectively) over their coincident events.




Fig. 23. Comparison plots of aerosol derived products. The panel on the left shows SAGE II aerosol surface area density in v7.0 minus v6.2 for all events after and including 1998. The panel on the right shows SAGE II aerosol effective radius in v7.0 minus v6.2 for all events after and including 1998.

latitude gridding or monthly means). The purpose of this fit was for the creation of climatologies, but has revealed some interesting data quality impacts between versions 6.2 and 7.0. Since the content of this paper focuses on the residuals of the fit rather than the fits themselves, we will only briefly outline the fitting process.

The fitting process begins by applying a modification of the Wang et al. (2002) filtering criteria to each event and then taking daily (zonal) means of collocated events. The following functional form is then regressed to all of the data:

$$\eta(\theta, t) = \Theta(\theta)T(t) \quad (6)$$




Fig. 24. Left – SAGE III and SAGE II mean water vapor mixing ratio profiles over all coincident events after shifting the water vapor channel center wavelength +12.0 nm. Right – SAGE III minus SAGE II water vapor comparison over coincident events after adjusting the water vapor channel in various ways: A – center wavelength +12.0 nm, FWHM +50 %; B – center wavelength +11.5 nm, FWHM +40 %, ozone effective cross-section in water vapor channel artificially decreased by 10 %; C – center wavelength +11.3 nm, FWHM +45 %, ozone cross-section database changed to Serdyuchenko et al. (2011); D – center wavelength +12.7 nm, FWHM +15 % (channel location prior to inclusion of MERRA in SAGE II v7.0).

where η is the concentration of the given species (i.e., O_3 , NO_2 , and H_2O), $\Theta(\theta)$ is the functional form of the latitudinal dependence, and $T(t)$ is the functional form of the time dependence. $\Theta(\theta)$ is simply a Fourier series (6 harmonics) with the constraint of zero derivative at the poles. $T(t)$ contains semiannual (3, 4, and 6 month terms), annual, quasi-biennial oscillation (QBO) (Singapore wind proxy at 30 hPa and 50 hPa), solar cycle (11 yr period terms), and equivalent effective stratospheric chlorine (EESC) (Newman et al., 2007) (for fits to O_3) terms as well as an additional piecewise term to account for any potential diurnal variation in a species. Both $\Theta(\theta)$ and $T(t)$ also contain a constant term, which collectively provide the constant for the fit. Residual analysis is performed to omit any outliers with large influence on the data and a correction is made for lag-1 autocorrelation (Reinsel et al., 1981). This fit is performed for each altitude. Some examples of the fits themselves are illustrated in Figs. 25 and 26. Two sets of plots are shown at low- and mid-latitudes to illustrate the robustness of the fit.

The absolute values of the residuals, hereafter simply stated as residuals, are examined throughout the fitting process. Figure 27 illustrates the mean residuals of the fit to O_3 prior to the lag-1 autocorrelation correction for versions 6.2 and 7.0 (total residuals). It also shows the same data averaged between 60° S and 60° N. There is a clear reduction in total residuals between 30 and 50 km from version 6.2 to version 7.0. The artifact in residuals in version 6.2 at 50 km is

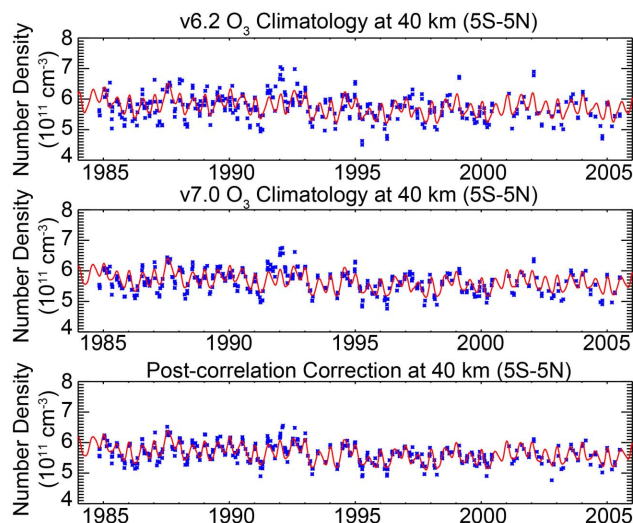


Fig. 25. Results of the new fitting routine on SAGE II O_3 data at 40 km. The data points include all data between 5° S and 5° N while the fit is plotted at the Equator. The first two plots show the data and fit for v6.2 and v7.0 prior to the lag-1 autocorrelation correction. The third plot shows the data and fit after the correction for v7.0, though the same plot for v6.2 is nearly identical. The residuals as a function of time are not shown as there was no discernible trend with time.

a result of smoothing applied above that altitude. Recall that version 6.2 employed a 2.5 km boxcar smoothing of O_3 data above ~ 50 km, while version 7.0 employs no smoothing to O_3 data. The correlated residuals, or those that are removed from the lag-1 autocorrelation correction, exhibit similar behavior as the total residuals (Fig. 28). The correlated residuals represent geophysical variability that is well sampled, but not well modeled. While the Singapore winds can be used as a reasonable proxy for the QBO in the tropics, it is still not exact and even less so at higher latitudes. The inability of the model to accurately represent the QBO can be seen as increased correlated (and total) residuals in regions where the model is less able to capture all of the geophysical variability that is attributable to the QBO. In addition, since the model attempts to fit all latitudes simultaneously, the deficiencies of using the proxy to model the QBO at higher latitudes will detract from the ability of the model to fit the QBO in the tropics. This effect is evident in Figs. 27 and 28. The uncorrelated residuals, or those that remain after the lag-1 autocorrelation correction, are, not surprisingly, similar between versions 6.2 and 7.0 (with the exception of the aforementioned smoothing) and reveal the extent of combined instrumental and geophysical noise in the data (Fig. 29). The uncertainties in the measurements (dashed lines) mirror the shape of this noise and are slightly (~ 1 %) lower than the observed residuals, indicating that the uncertainty estimates are reasonable and that the zonally averaged, unresolved geophysical variability




Fig. 26. Results of the new fitting routine on SAGE II O₃ data at 40 km. The data points include all data between 40° N and 50° N while the fit is plotted at 45° N. The first two plots show the data and fit for v6.2 and v7.0 prior to the lag-1 autocorrelation correction. The third plot shows the data and fit after the correction for v7.0, though the same plot for v6.2 is nearly identical. The residuals as a function of time are not shown as there was no discernible trend with time.




Fig. 27. Results of the fitting algorithm to O₃ showing altitude and latitude dependence of the mean absolute residual prior to the lag-1 autocorrelation correction for both v6.2 and v7.0. The mean absolute residual at each altitude is also shown.

is on the order of 1 %. The uncorrelated residuals are slightly higher in the tropics near 30 km, which is a result of geophysical variability that is not well sampled as opposed to a localized region of increased noise. All told, version 7.0 data is more consistent, with less correlated noise and smaller total residuals than version 6.2, making the data more robust to work with for time series analysis.

Figure 30 shows the same (non-latitude dependent) plots as O₃ for sunset NO₂. There is again a clear overall decrease in the total residuals in the mid-stratosphere. The artifact in version 6.2 at 40 km is a result of the transition to a “no

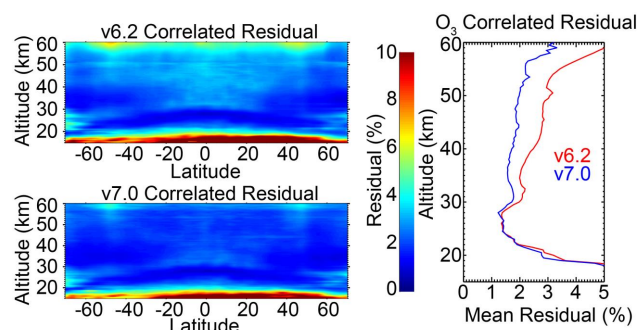


Fig. 28. Results of the fitting algorithm to O₃ showing altitude and latitude dependence of the mean absolute residual as a result of the lag-1 autocorrelation for both v6.2 and v7.0. The mean absolute residual at each altitude is also shown.

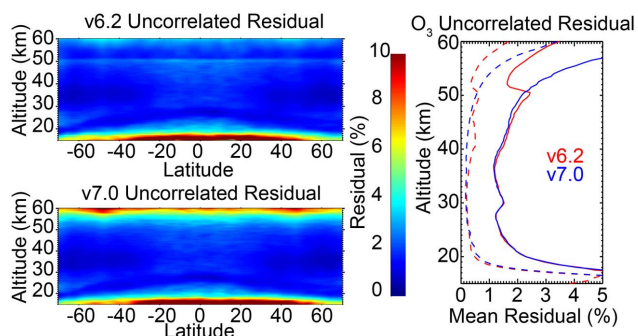


Fig. 29. Results of the fitting algorithm to O₃ showing altitude and latitude dependence of the mean absolute residual after the lag-1 autocorrelation correction for both v6.2 and v7.0. The mean absolute residual at each altitude is also shown, with the mean uncertainty in the measurements shown as dashed lines.

aerosol” retrieval method above that altitude in addition to additional smoothing. While the correlated residuals have decreased from version 6.2 to version 7.0 as expected, they exhibit different behavior above 40 km in version 7.0, which is a result of the previously discussed high bias in sunset NO₂. Analysis of the uncorrelated residuals for sunsets shows behavior consistent with that of O₃ (i.e., version 6.2 and version 7.0 are comparable, with mean uncertainties that mimic their shapes at a smaller level and an implied geophysical variability of ~ 1 %) above 25 km and below the high bias in the version 7.0 data. At altitudes below 25 km, version 7.0 has more noise than in version 6.2 as well as having large uncertainties. The cause of this behavior is a matter of further research. A look at sunrise NO₂ (Fig. 31) reveals similar, although much less consistent, results as sunset NO₂. Sunrise NO₂ shows consistent measurements only near the peak, but it is still a research product.

Figure 32 shows the fitting results for water vapor. While the total residuals have decreased overall between version 6.2 and version 7.0, especially at higher altitudes thanks to changes in the retrieval, the behavior of the correlated

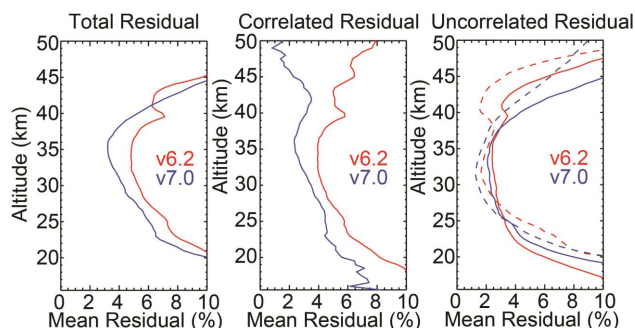


Fig. 30. Results of the fitting algorithm to sunset NO_2 , showing the mean absolute residual at each altitude for both v6.2 and v7.0 at the various stages of the fitting process. The plot of uncorrelated residuals also shows the mean uncertainty in the measurements as dashed lines.

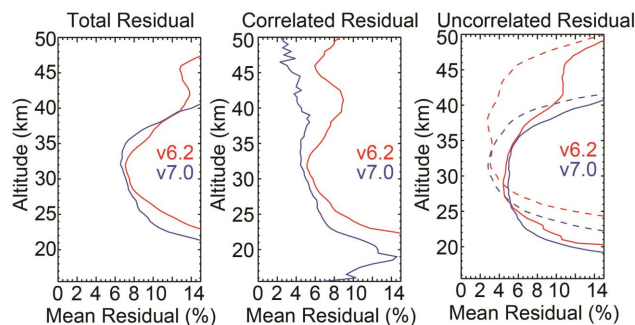


Fig. 31. Results of the fitting algorithm to sunrise NO_2 , showing the mean absolute residual at each altitude for both v6.2 and v7.0 at the various stages of the fitting process. The plot of uncorrelated residuals also shows the mean uncertainty in the measurements as dashed lines.

residuals is not consistent with other species. Several possible causes exist to explain this behavior, though we believe the most likely explanation is a lingering interference from other species in the water vapor product. The uncertainties in the water vapor measurements are much too high when compared to the uncorrelated residuals at lower altitudes (there is a transition to additional smoothing above 25 km), which may be the result of the inclusion of an aerosol clearing uncertainty into the water vapor uncertainty. The aerosol uncertainty manifests itself as a bias rather than increased random noise. As previously mentioned, we are still not satisfied with the result and plan on addressing these issues in a future revision.

With the exception of some unusual behavior in water vapor, the total residuals of retrieved species have decreased from version 6.2 to version 7.0, partly as a result of improvements to the algorithm and partly as a result of decreased correlated residuals. Since the correlated residuals are the result of the lag-1 autocorrelation, and the data has been consolidated into daily means, the correlated residuals represent a combination of daily geophysical variability that is not

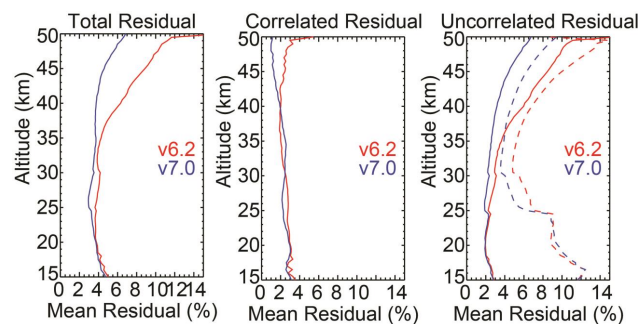


Fig. 32. Results of the fitting algorithm to H_2O , showing the mean absolute residual at each altitude for both v6.2 and v7.0 at the various stages of the fitting process. The plot of uncorrelated residuals also shows the mean uncertainty in the measurements as dashed lines.




Fig. 33. Results of the fitting algorithm to neutral density, showing the mean absolute residual at each altitude for both v6.2 and v7.0 at the various stages of the fitting process. The plot of uncorrelated residuals also shows the mean uncertainty in the measurements as dashed lines.

captured in the model and noise in the data that is correlated from day to day. Since the algorithm itself cannot produce correlated noise (without serious errors), any correlated noise in the data must be a result of correlated noise in the input data, namely the ephemeris and meteorological data. While the possibility of correlated noise in the ephemeris data itself exists (Buglia, 1989), the manner in which the algorithm processes ephemeris data and the fact that ephemeris data is highly correlated from day to day can affect the correlated noise in the processed ephemeris data. As such, version 6.2 did have a source of correlated ephemeris noise. The altitude registration offset caused by improper ephemeris calculations (outlined in Sect. 2.1 and shown in Fig. 2) would be interpreted as a source of correlated noise in the processed ephemeris data. Since the solar ephemeris calculations in version 7.0 are more accurate, this correlated noise in the processed ephemeris data has been reduced. The meteorological input data also have some amount of correlated noise. Figure 33 illustrates the same residual analysis for the neutral density that comes from the meteorological data. While the

overall level of correlated noise is small, there is a clear reduction in correlated noise above ~ 35 km in version 7.0. Recall from Sect. 2.2 that, in version 6.2, operational model data provided by NCEP were used between 10 mbar and 0.4 mbar (~ 30 – 50 km) and have been shown to be potentially inconsistent and problematic. This explains the separation of the correlated residuals between versions 6.2 and 7.0 and the subsequent convergence as both algorithms eventually return to the use of GRAM-95 data at higher altitudes. The reduction of both of these sources of correlated input data result in the subsequent reduction of correlated noise in the retrieved species.

6 Conclusions

Version 7.0 of the SAGE II processing algorithm has been detailed and discussed and the changes made between version 6.2 and version 7.0 have been explained. Some of the changes from version 6.2 to version 7.0 that have had the greatest impacts on the data include corrections to solar ephemeris processing to remove a quasi-random altitude registration offset, migration from NCEP ancillary meteorological data to MERRA, adaptation of updated O_3 and NO_2 temperature-dependent absorption cross-sections, incorporation of transmission processing algorithms closer to those used in SAGE III version 4.0, redetermination of the water vapor channel filter spectral location due to spectral drift, and replacing Twomey–Chahine inversion in favor of a simple onion peeling inversion technique.

Ozone has always been and continues to be a high quality data product. As a result of changes made in version 7.0, the mean value of ozone has decreased by 1.5 %, but is now more consistent as a function of altitude when compared with other instruments. Ozone is also now more robust for use in time-series analysis, with smaller total residuals and decreased correlated noise in the data. The continued use of SAGE II ozone data is encouraged, though care should be used when using the data either below 20 km or above 50 km due to increased noise in the data.

NO_2 continues to be a good product for study, though use of only sunset events is recommended as sunrise NO_2 , while better in version 7.0 than version 6.2, is still somewhat of a research product. Version 7.0 sunset NO_2 concentrations have decreased by 5–10 % and are in better agreement with SAGE III than were those of version 6.2. Much like O_3 , sunset NO_2 is now easier to work with for time-series analysis, with smaller total residuals and decreased correlated noise in the data. Care should be taken when using the data below 25 km or above 40 km due to noise. In addition, sunset NO_2 values above 40 km are biased high and an algorithm to better retrieve NO_2 at these altitudes is in development.

Aerosol has also been a cornerstone SAGE II data products and continues to be so in version 7.0. Changes to spectroscopy in version 7.0 have resulted in large changes in

aerosol extinction in the 525 and 452 nm channels. While 1020 nm aerosol has remained consistent, 525 nm aerosol is now in better agreement with SAGE III, though 452 nm aerosol is now much smaller than before. Use of aerosol in the 385 nm channel continues to be discouraged due to unknown issues with that channel. The aerosol derived products have been updated, with surface area density increasing by 50 % during clean periods and effective radius decreasing by 10 % compared to version 6.2.

Water vapor continues to be a work in progress. Many changes have been made to the version 7.0 algorithm to improve the quality of the water vapor data product, though several sources of uncertainty remain that require additional work. Most notably is the presence of lingering signatures from other interfering species. While potential solutions are currently being investigated, we did not want to hold back the release of the significant improvements made to the O_3 and aerosol products. That having been said, the SAGE II version 7.0 water vapor data product is now in much better agreement with more modern data sets and shows significant improvement above 40 km compared to previous versions, though we still strongly caution the use of this data for trend studies.

Acknowledgements. Refinement of algorithms for the SAGE series of missions is supported by the Earth Science Division of NASA's Science Mission Directorate.

Edited by: D. Loyola

References

- Auer, L. H. and Standish, E. M.: Astronomical Refraction: Computational Method for All Zenith Angles, *Astron. J.*, 119, 2472–2472, doi:10.1086/301325, 2000.
- Bogumil, K., Orphal, J., Homann, T., Voigt, S., Spietz, P., Fleischmann, O., Vogel, A., Hartmann, M., Bovensmann, H., Frerik, J., and Burrows, J.: Measurements of molecular absorption spectra with the SCIAMACHY pre-flight model: Instrument characterization and reference data for atmospheric remote sensing in the 230–2380 nm region, *J. Photoch. Photobio. A*, 157, 167–184, doi:10.1016/S1010-6030(03)00062-5, 2003.
- Bucholtz, A.: Rayleigh-scattering calculations for the terrestrial atmosphere, *Appl. Optics*, 34, 2765–2773, doi:10.1364/AO.34.002765, 1995.
- Buglia, J. J.: Compilation of methods in orbital mechanics and solar geometry, NASA Reference Publication 1204, 1988.
- Buglia, J. J.: Effect of ephemeris errors on the accuracy of the computation of the tangent point altitude of a solar scanning ray as measured by the SAGE I and II instruments, NASA Technical Paper 2866, 1989.
- Burton, S. P. and Thomason, L. W.: Polar mesospheric clouds in SAGE II Version 6.0 data, Proceedings of the Quadrennial Ozone Symposium, Sapporo, Japan, International Ozone Commission, 325–326, 2000.
- Burton, S. P., Thomason, L. W., and Zawodny, J. M.: Technical Note: Time-dependent limb-darkening calibration for so-

- lar occultation instruments, *Atmos. Chem. Phys.*, 10, 1–8, doi:10.5194/acp-10-1-2010, 2010.
- Chahine, M. T.: A general relaxation method for inverse solution of the full radiative transfer equation, *J. Atmos. Sci.*, 29, 741–747, 1972.
- Chehade, W., Gür, B., Spietz, P., Gorshchev, V., Serdyuchenko, A., Burrows, J. P., and Weber, M.: Temperature dependent ozone absorption cross section spectra measured with the GOME-2 FM3 spectrometer and first application in satellite retrievals, *Atmos. Meas. Tech.*, 6, 1623–1632, doi:10.5194/amt-6-1623-2013, 2013.
- Chu, D. A. and Cunnold, D. M.: Mesospheric ozone measurements by SAGE II, NASA GSFC, Ozone in the Troposphere and Stratosphere, Part 2, 895–898 (SEE N95-11006 01-47), 1994.
- Chu, W. P.: Calculations of atmospheric refraction for spacecraft remote-sensing applications, *Appl. Optics*, 22, 721–725, doi:10.1364/AO.22.000721, 1983.
- Chu, W. P., McCormick, M. P., Lenoble, J., Broniez, C., and Pruvost, P.: SAGE II inversion algorithm, *J. Geophys. Res.*, 94, 8339–8351, doi:10.1029/JD094iD06p08447, 1989.
- Chu, W. P., Chiou, E. W., Larsen, J. C., Thomason, L. W., Rind, D., Buglia, J. J., Oltmans, S., McCormick, M. P., and McMaster, L. R.: Algorithms and sensitivity analyses for SAGE II water vapor retrieval, *J. Geophys. Res.*, 98, 4857–4866, 1993.
- Cunnold, D. M., Chu, W. P., Barnes, R. A., McCormick, M. P., and Veiga, R. E.: Validation of SAGE II Ozone Measurements, *J. Geophys. Res.*, 94, 8447–8460, 1989.
- Finger, F. G., Gelman, M. E., Wild, J. D., Chanin, M. L., Hauchecorne, A., and Miller, A. J.: Evaluation of NCEP Upper-stratospheric temperature analyses using rocketsondes and lidar data, *B. Am. Meteorol. Soc.*, 74, 789–799, 1993.
- Gordley, L. L. and Russell, J. M.: A fast and accurate radiance algorithm for applications to inversion of limb measurements, *Adv. Space Res.-Series*, edited by: Deepak, A., Academic Press, 591–607, 1980.
- Johnson, D. L., Justus, C. G., Jeffries III, W. R., and Yung, S. P.: The NASA/MSFC Global Reference Atmospheric Model – 1995 Version (GRAM-95), NASA Technical Memorandum 4715, 1995.
- Johnston, P. V. and McKenzie, R. L.: Long-path absorption measurements of tropospheric NO₂ in rural New Zealand, *Geophys. Res. Lett.*, 11, 69–72, 1984.
- Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., Mo, K. C., Ropelewski, C., Wang, J., Leetmaa, A., Reynolds, R., Jenne, R., and Joseph, D.: The NCEP/NCAR 40-year reanalysis project, *B. Am. Meteorol. Soc.*, 77, 437–470, 1996.
- Kurucz, R. L., Furenlid, I., Brault, J., and Testerman, L.: Solar Flux atlas from 296 to 1300 nm, National Solar Observatory Atlas No. 1, NAO, Sunspot, NM, 1984.
- Lambert, A., Reade, W. G., Livesey, N. J., Santee, M. L., Manney, G. L., Froidevaux, L., Wu, D. L., Schwartz, M. J., Pumphrey, H. C., Jimenez, C., Nedoluha, G. E., Cofield, R. E., Cuddy, D. T., Daffer, W. H., Drouin, B. J., Fuller, R. A., Jarnot, R. F., Knosp, B. W., Pickett, H. M., Perun, V. S., Snyder, W. V., Stek, P. C., Thurstans, R. P., Wagner, P. A., Water, J. W., Jucks, K. W., Toon, G. C., Stachnik, R. A., Bernath, P. F., Boone, C. D., Walker, K. A., Urban, J., Murtagh, D., Elkins, J. W., and Atlas, E.: Validation of the Aura Microwave Limb Sounder middle atmosphere water vapor and nitrous oxide measurements, *J. Geophys. Res.*, 112, D24S36, doi:10.1029/2007JD008724, 2007.
- Lemoine, F. G., Kenyon, S. C., Factor, J. K., Trimmer, R. G., Pavlis, N. K., Chinn, D. S., Cox, C. M., Klosko, S. M., Luthcke, S. B., Torrence, M. H., Wang, Y. M., Williamson, R. G., Pavlis, E. C., Rapp, R. H., and Olson, T. R.: The Development of the Joint NASA GSFC and NIMA Geopotential Model EGM96, NASA Technical Publication 1998–206861, 1998.
- Mauldin, L. E., Zaun, N. H., McCormick, M. P., Guy, J. H., and Vaughn, W. R.: Stratospheric Aerosol and Gas Experiment II Instrument: A Functional Description, *Opt. Eng.*, 24, 307–312, doi:10.1117/12.7973473, 1985.
- Mlawer, E. J., Clough, S. A., Brown, P. D., Stephen, T. M., Landry, J. C., Goldman, A., and Murcray, F. J.: Observed atmospheric collision-induced absorption in near-infrared oxygen bands, *J. Geophys. Res.*, 103, 3859–3863, doi:10.1029/97JD03141, 1998.
- Newman, P. A., Daniel, J. S., Waugh, D. W., and Nash, E. R.: A new formulation of equivalent effective stratospheric chlorine (EESC), *Atmos. Chem. Phys.*, 7, 4537–4552, doi:10.5194/acp-7-4537-2007, 2007.
- Newnham, D. A. and Ballard, J.: Visible absorption cross sections and integrated absorption intensities of molecular oxygen (O₂ and O₄), *J. Geophys. Res.*, 103, 28801–28815, doi:10.1029/98JD02799, 1998.
- NIMA Technical Report TR8350.2: Department of Defense World Geodetic System 1984, Its Definition and Relationships With Local Geodetic Systems, 3rd Edn., 4 July 1997.
- Noerdlinger, P. D.: Theoretical Basis of the SDP Toolkit Geolocation Package for the ECS Project, Hughes Applied Information Systems Technical Publication 445-TP-002-002, 201 pp., 1995.
- Pitts, M. C., Thomason, L. W., Zawodny, J. M., Wenny, B. N., Livingston, J. M., Russell, P. B., Yee, J.-H., Swartz, W. H., and Shetter, R. E.: Ozone observations by the Gas and Aerosol Measurement Sensor during SOLVE II, *Atmos. Chem. Phys.*, 6, 2695–2709, doi:10.5194/acp-6-2695-2006, 2006.
- Reinsel, G., Tiao, G. C., Wang, M. N., Lewis, R., and Nychka, D.: Statistical analysis of stratospheric ozone data for the detection of trends, *Atmos. Environ.*, 15, 1569–1577, doi:10.1016/0004-6981(81)90140-2, 1981.
- Rienecker, M. M., Suarez, M. J., Todling, R., Bacmeister, J., Takacs, L., Liu, H.-C., Gu, W., Sienkiewicz, M., Koster, R. D., Gelaro, R., Stajner, I., and Nielsen, J. E.: The GEOS-5 Data Assimilation System – Documentation of Versions 5.0.1, 5.1.0, and 5.2.0, Technical Report Series on Global Modeling and Data Assimilation, Vol. 27, 2008.
- Rienecker, M. M., Suarez, M. J., Gelaro, R., Todling, R., Bacmeister, J., Liu, E., Bosilovich, M. G., Schubert, S. D., Takacs, L., Kim, G.-K., Bloom, S., Chen, J., Collins, D., Conaty, A., da Silva, A., Gu, W., Joiner, J., Koster, R. D., Lucchesi, R., Molod, A., Owens, T., Pawson, S., Pegion, P., Redder, C. R., Reichle, R., Robertson, F. R., Ruddick, A. G., Sienkiewicz, M., and Woollen, J.: MERRA – NASA’s Modern-Era Retrospective Analysis for Research and Applications, *J. Climate*, 24, 3624–3648, doi:10.1175/JCLI-D-11-00015.1, 2011.
- Rothman, L. S., Jacquemart, D., Barbe, A., Benner, D. C., Birk, M., Brown, L. R., Carleer, M. R., Chackerian Jr., C., Chance, K., Coudert, L. H., Dana, V., Devi, V. M., Flaud, J.-M., Gamache, R. R., Goldman, A., Hartmann, J.-M., Jucks, K. W., Maki, A. G., Mandin, J.-Y., Massie, S. T., Orphal, J., Perrin, A., Rinsland, C.

- P., Smith, M. A. H., Tennyson, J., Tolchenov, R. N., Toth, R. A., Auwera, J. V., Varanasi, P., and Wagner, G.: The HITRAN 2004 molecular spectroscopic database, *J. Quant. Spectrosc. Ra.*, 96, 139–204, doi:10.1016/J.QSRT.2004.10.008, 2004.
- Rothman, L. S., Gordon, I. E., Barbe, A., Benner, D. C., Bernath, P. F., Birk, M., Boudon, V., Brown, L. R., Campargue, A., Champion, J.-P., Chance, K., Coudert, L. H., Dana, V., Devi, V. M., Fally, S., Flaud, J.-M., Gamache, R. R., Goldman, A., Jacquemart, D., Kleiner, I., Lacome, N., Lafferty, W. J., Mandin, J.-Y., Massie, S. T., Mikhailenko, S. N., Miller, C. E., Moazzen-Ahmadi, N., Naumenko, O. V., Nikitin, A. V., Orphal, J., Perevalov, V. I., Perrin, A., Predoi-Cross, A., Rinsland, C. P., Rotger, M., Simeckova, M., Smith, M. A. H., Sung, K., Tashkun, S. A., Tennyson, J., Toth, R. A., Vandaele, A. C., and Auwera, J. V.: The HITRAN 2008 molecular spectroscopic database, *J. Quant. Spectrosc. Ra.*, 110, 533–572, doi:10.1016/J.QSRT.2009.02.013, 2009.
- Russell III, J. M., Gordley, L. L., Park, J. H., Drayson, S. R., Hesketh, D. H., Cicerone, R. J., Tuck, A. F., Frederick, J. E., Harries, J. E., and Crutzen, P.: The Halogen Occultation Experiment, *J. Geophys. Res.*, 98, 10777–10797, doi:10.1029/93JD00799, 1993.
- Serdyuchenko, A., Gorshelev, V., Weber, M., and Burrows, J. P.: New broadband high-resolution ozone absorption cross-sections, *Spectroscopy Europe/Asia*, 23, 14–17, 2011.
- Shettle, E. P. and Anderson, S.: New visible and near IR ozone absorption cross-sections for MODTRAN, Proceedings of the 17th Annual Review Conference on Atmospheric Transmission Models, 8–9 June 1994, edited by: Anderson, G. P., Picard, R. H., and Chetwynd, J. H., PL-TR-95-2060, Phillips Laboratory, Hanscom AFB, MA, 335–345, 1995.
- SPARC/IOC/GAW (Stratospheric Processes and their Role in Climate/International Ozone Commission/Global Atmosphere Watch): Assessment of Trends in the Vertical Distribution of Ozone, SPARC Rep. 1, WMO-Ozone Res. Monit. Proj. Rep. 43, edited by: Harris, N., Hudson, R., and Phillips, C., World Meteorological Organization, Geneva, 1998.
- Thomason, L. W., Poole, L. R., and Deshler, T. R.: A global climatology of stratospheric aerosol surface area density as deduced from SAGE II: 1984–1994, *J. Geophys. Res.*, 102, 8967–8976, doi:10.1029/96JD02962, 1997.
- Thomason, L. W., Burton, S. P., Iyer, N., Zawodny, J. M., and Anderson, J.: A revised water vapor product for the Stratospheric Aerosol and Gas Experiment (SAGE) II version 6.2 data set, *J. Geophys. Res.*, 109, D06312, doi:10.1029/2003JD004465, 2004.
- Thomason, L. W., Burton, S. P., Luo, B.-P., and Peter, T.: SAGE II measurements of stratospheric aerosol properties at non-volcanic levels, *Atmos. Chem. Phys.*, 8, 983–995, doi:10.5194/acp-8-983-2008, 2008.
- Thomason, L. W., Moore, J. R., Pitts, M. C., Zawodny, J. M., and Chiou, E. W.: An evaluation of the SAGE III version 4 aerosol extinction coefficient and water vapor data products, *Atmos. Chem. Phys.*, 10, 2159–2173, doi:10.5194/acp-10-2159-2010, 2010.
- Twomey, S.: Comparison of constrained linear inversion and an iterative nonlinear algorithm applied to the indirect estimation of particle size distributions, *J. Comput. Phys.*, 18, 188–198, 1975.
- Wang, H.-J., Cunnold, D. M., Thomason, L. W., Zawodny, J. M., and Bodeker, G. E.: Assessment of SAGE version 6.1 ozone data quality, *J. Geophys. Res.*, 107, 4691, doi:10.1029/2002JD002418, 2002.
- Wang, H.-J., Cunnold, D. M., Trepte, C., Thomason, L. W., and Zawodny, J. M.: SAGE III solar ozone measurements: Initial results, *Geophys. Res. Lett.*, 33, L03805, doi:10.1029/2005GL025099, 2006.
- WMO (World Meteorological Organization): Scientific Assessment of Ozone Depletion: 2010, Global Ozone Research and Monitoring Project-Report No. 52, 516 pp., Geneva, Switzerland, 2011.