



Contents lists available at ScienceDirect

Journal of Quantitative Spectroscopy & Radiative Transfer

journal homepage: www.elsevier.com/locate/jqsrt

Absorption lookup tables in the radiative transfer model ARTS

Stefan A. Buehler^{a,*}, Patrick Eriksson^b, Oliver Lemke^a^a Division of Space Technology, Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, Box 812, SE-98128 Kiruna, Sweden^b Department of Earth and Space Sciences, Chalmers University of Technology, SE-41296 Gothenburg, Sweden

ARTICLE INFO

Available online 16 March 2011

Keywords:

Radiative transfer models

Absorption

Lookup table

ARTS

AMSU

HIRS

Odin

ABSTRACT

We describe the lookup table approach that is used to store pre-calculated absorption data in the radiative transfer model ARTS. The table stores absorption cross sections as a function of frequency, pressure, temperature, and the water vapor volume mixing ratio, where the last dimension is only included for those gas species that require it. The table is used together with an extraction strategy, which uses polynomial interpolation, with recommended interpolation orders between five and seven. We also derived recommended default settings for grid spacings and interpolation orders, and verified that the approach gives very accurate results with these default settings. The tested instrument setups were for AMSU-B, HIRS, and Odin, three well-known satellite remote sensing instruments covering a wide range of frequencies and viewing geometries. Errors introduced by the lookup table were found to be always below a few millikelvin, in terms of the simulated brightness temperature.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

The atmospheric radiative transfer simulator (ARTS) is a public domain radiative transfer model for thermal radiation in planetary atmospheres [1,2]. It is freely available on the Internet at <http://www.sat.ltu.se/arts>. The model is applicable to frequencies from the microwave to the thermal infrared. Areas of application include the simulation of remote measurements and the calculation of Jacobians for remote measurement inversion [3], but also the accurate simulation of broad band radiation fluxes [4,5].

In contrast to fast parametric models like RTTOV [6], ARTS is a physical model. With that we mean, firstly, that it calculates absorption coefficients line-by-line from spectroscopic catalogues (plus continua), whereas for

example RTTOV computes optical depths as a linear combination of profile dependent predictors [7]. In both cases absorption and optical depth will depend on atmospheric pressure, temperature, and trace gas concentrations, but in the RTTOV case the optical depth will depend also on viewing angle, which is for ARTS handled by the radiative transfer scheme.

Secondly, ARTS numerically integrates the monochromatic radiative transfer equation (RTE) for a discrete set of frequencies, whereas for example RTTOV uses the polychromatic form of the RTE, which operates on channel average transmissions [7]. For the clear-sky case, solving the RTE involves just an integral along the line of sight. If scatterers, such as cloud particles, are present, the solution is more complicated (see below).

The fact that ARTS is a physical model has the advantage that it is accurate, especially for atmospheric situations that are poorly represented in the training data, where parametric models may have increased errors. An example for such a situation, for the specific case of RTTOV-7 and the AMSU-B sensor, is discussed in [8]. It is reasonable to assume that even other fast models have

* Corresponding author. Tel.: +4698079177; fax: +4698079050.

E-mail addresses: sbuehler@ltu.se (S.A. Buehler),patrick.eriksson@chalmers.se (P. Eriksson), olemke@ltu.se (O. Lemke).URLS: <http://www.sat.ltu.se> (S.A. Buehler), <http://www.chalmers.se/rss> (P. Eriksson), <http://www.sat.ltu.se> (O. Lemke).

larger errors for ‘exotic’ atmospheric situations that are not sufficiently represented in the training data.

However, the high accuracy of ARTS comes at the price of high computational cost. In particular, the calculation of line-by-line absorption coefficients is costly, because it can involve the summation of contributions from thousands or even tens of thousands of spectral lines for each calculation frequency. There are two important cases where this line-by-line absorption calculation has to be done many times over in a straightforward implementation. Case one is the simulation of spectra for a batch of atmospheric states, using the same sensor setup. This situation occurs frequently, both when processing real sensor data, and when generating training data sets. Case two is the simulation with a 2D or 3D atmosphere, where each point in the 2D or 3D grid requires different absorption coefficients, depending on its pressure, temperature, and trace gas concentrations. In such cases, where line-by-line absorption would have to be calculated repeatedly, it is efficient to pre-calculate absorption coefficients, and store them in a lookup table. While not making ARTS as fast as a true ‘fast’ model, this approach makes it fast enough to open a new applicability range, such as the generation of training data.

Another important case, where a lookup table can significantly improve the computational efficiency of ARTS, is radiative transfer with hydrometeor scattering. We need such calculations for example in the context of submillimeter-wave remote sensing of cloud ice [9,10]. ARTS includes two solvers for simulations with scattering, one Monte Carlo solver [11,12], and one iterative discrete ordinate solver [13,14]. Although it is technically possible in ARTS to run these solvers and explicitly calculate line-by-line absorption ‘on the fly’, the calculation speed with this setup will be too slow for most applications. Hence, in practice both solvers are normally be used together with an absorption lookup table.

With absorption lookup table we mean that absorption is calculated line-by-line on predefined grids of pressure, temperature, frequency, and possibly water vapor amount, and then stored in a table for later use. A second important element of any absorption lookup table implementation is the extraction strategy, which defines how actual absorption for specific atmospheric conditions can be extracted from the table. This strategy usually involves interpolation.

This article describes the implementation of such a lookup table strategy in ARTS, documents its validation, and gives recommendations for its optimal use. The general approach is not new, and we therefore give an overview of different documented earlier implementations in Section 2. Section 3 then describes the lookup table format and extraction algorithm. Section 4 describes its testing and validation, and contains some usage recommendations. Finally, Section 5 contains a summary and the conclusions, including a brief discussion of the advantages and disadvantages of the chosen approach.

2. Earlier implementations

The aim of this section is to give a compact overview of earlier absorption lookup table implementations. As a

complement to the text, Table 1 lists all discussed models and their main properties. For the sake of brevity we do not thoroughly define all terms here. The reader that is not familiar with this subject may find it useful to jump ahead and read Section 3 first, where the terms are defined, and the encountered issues are discussed in more detail.

The earliest reported implementation that we are aware of is that of [16] in 1981. Their table treated the frequency, pressure, and temperature dependence of the absorption for different absorbers. No special treatment was given to water vapor, so water vapor self-broadening was ignored. As in many later implementations (but not all), the temperature grid was different for different pressures, in order to avoid calculating and storing absorption for p/T combinations that do not occur in the atmosphere. We also use this approach. Furthermore, to save some storage space, pressure/frequency combinations where transmission is close to one were omitted.

Aoki [17] in 1988 used a quite similar approach, but with some modifications. He stores (and interpolates) the square-root of the absorption, in order to improve interpolation accuracy, and he uses a quadratic form for the temperature interpolation. To save space, he uses different non-equidistant frequency grids for different pressures.

In 1995, [18] developed his own absorption lookup table implementation, seemingly unaware of both [16] and [17]. He stores and interpolates the logarithm of the absorption coefficient, on a regular frequency, pressure, and temperature grid. (So, the temperature grid is the same for all pressures.) His extraction strategy is a two-dimensional cubic polynomial interpolation in pressure ($\log(p/p_0)$) and temperature (T/T_0). In this respect there is a similarity with our work, since we also use polynomial interpolation, albeit usually with higher orders. Turner made no attempts to reduce the size of his table or to compress it.

Only two years later, in 1998, [19] published a landmark paper on the subject. Their interpolation scheme was again slightly different, and more importantly they introduced a method of compression using singular value decomposition (SVD) to reduce the table memory requirements, which two other implementations since then have adopted. Our own implementation so far does not use this method, but it could be added relatively easily. [19] were also the first to introduce a special treatment of water vapor, to account for the self-broadening effect in the lower troposphere. We also do include this effect.

In 2001 and 2002, [20] and [21] published their respective implementations. Both use the SVD compression technique, but combine it with quite different interpolation schemes (see Table 1 for details). Another difference is that [20] do include water vapor self-broadening, whereas [21] do not. Since then, two more implementations have been published that we are aware of, [15,22], both without compression, and each using different combinations of grids and interpolation (see Table 1).

It is difficult to draw any general conclusions from these earlier implementations. Storage method, grids, and interpolation schemes form a unit that can only be judged

Table 1

An overview of published absorption lookup table implementations. Notes: (1) T -pert in the 'dimensions' column means that a different temperature grid is used for each pressure value, to avoid p/T combinations that do not occur in the atmosphere. (2) τ is the layer optical depth, the integral of α across the layer, κ is the absorption cross section, defined as α divided by absorber number density (see Eq. (2)). (3) Column 'H₂O dep.' states to what extent there is a special treatment for H₂O, 'self and O₂' means that both H₂O and O₂ absorption are stored for different H₂O values. (4) Those implementations that do treat H₂O in a special way will also do some interpolation in that dimension, to keep the table here concise, the interpolation scheme for this is not listed. (5) Ref. [15] was developed for the atmosphere of Venus, all other implementations are for the atmosphere of Earth.

Authors	Year	Dimensions	H ₂ O dep.	Quantity	Interpolation	Compression
Scott and Chedin [16]	1981	Absorber, v, p, T -pert	None	τ	T : nearest neighbour p : none v : none	omit layers/frequencies where transmittance close to 1.
Aoki [17]	1988	Absorber, v, p, T -pert	None	$\sqrt{\alpha}$	T : quadratic form p : linear v : linear	variable v grid, depending on p .
Turner [18]	1995	Absorber, v, p, T	None	$\log(\alpha)$	2D cubic polynomial interpolation in $\log(p/p_0)$ and T/T_0 v : none	none
Strow et al. [19]	1998	Absorber, v, p, T -pert	Self	τ or $\sqrt[3]{\tau}$	T : 5th order polynomial p : none v : none	SVD compression
Mitsel et al. [20]	2001	Absorber, v, p, T -pert	Self	α	Square interpolation in T, p, v	SVD compression, plus variable v grid
Dudhia et al. [21]	2002	Absorber, v, p, T	None	$\ln(\alpha)$	T : linear p : linear in $\ln(p)$ v : linear	SVD compression, plus different T, p, v grids for each absorber
Koukouli et al. [15]	2005	Absorber, v, p, T	None	α	Unclear	none
Clough et al. [22]	2006	Absorber, v, p, T -pert	Self	α	T : unclear p : none v : none	none
This work	2011	Absorber, v, p, T -pert	Self, O ₂	κ	T : polynomial (default order 7) p : polynomial (default order 5) v : none	none

as a whole, and in the context of the application it was developed for. One can definitely say that a special treatment of water vapor is necessary if the model is used for the lower troposphere. Furthermore, based on our own experiments described later, one can say that higher order interpolation schemes tend to give more accurate results for the same grid spacing, as expected. Lastly, one can say that compression is necessary where total table size is a limiting factor. But the need for this has somewhat decreased in recent years, due to the continuing growth of available computer memory.

The next section describes our own absorption lookup table implementation in some detail. We do not claim it to be superior to earlier implementations. However, it does have the distinct advantages of being easy to use, quite flexible, and freely available. At least the last point probably distinguishes it from most of the other implementations.

3. Method

3.1. Absorption coefficients and absorption cross sections

Although ARTS handles scattering atmospheres [13], for simplicity, we will use the clear-sky radiative transfer equation as the basis of discussion. We also assume local thermal equilibrium, without discussing this further here.

As described for example in [23], the RTE under these assumptions is

$$\frac{dI}{ds} = -\alpha I + \alpha B(T), \quad (1)$$

where I is the specific intensity, defined as the flux of energy in a given direction per second per unit frequency interval per unit solid angle per unit area. The variable s is the distance along the propagation path, and α is the absorption coefficient in units of 1/length. The function $B(T)$ is the Planck function, and T is temperature.

Eq. (1) has been written without indicating variable dependences, in order to keep it simple and compact. But it should be noted that I is actually $I(s)$, and T is actually $T(s)$. Furthermore, the absorption coefficient α is a monochromatic quantity, so we have to calculate (and store) it for any frequency ν that we use in the simulation. So α is actually

$$\alpha(\nu, p(s), T(s), x_1(s), \dots, x_N(s)),$$

where ν is frequency, $p(s)$ is total atmospheric pressure, $T(s)$ is temperature, and the $x_i(s)$ are trace gas volume mixing ratios (VMRs).

Lambert's law of extinction (see for example [23]) states that the decrease in intensity along the propagation path is proportional to the intensity itself (first term in Eq. (1)), and proportional to the amount of matter in the

path. Indirectly, this implies that for a mixture of different gases the total absorption can be written as a sum of the partial absorption for each gas. We can thus rewrite α as

$$\alpha = \sum_i \alpha_i = \sum_i \kappa_i n_i, \quad (2)$$

where the index i denotes the different gas species, the n_i are the number densities of the different species, and the κ_i are the absorption cross sections of the different species. The unit of κ_i is that of an area (length^2), hence the name. Like α , κ_i is actually

$$\kappa_i(v, p(s), T(s), x_1(s), \dots, x_N(s)).$$

As one can see from Eq. (2), a significant pressure dependence of the α_i comes from the n_i . (If one assumes a constant VMR of species i , then n_i is proportional to the total pressure according to the ideal gas law.)

When constructing an absorption lookup table, it is preferable to store the κ_i , rather than the α_i . The reason for this is that one has to interpolate in pressure. The less the interpolated quantity varies with pressure, the smaller the interpolation errors.

One then has to think about the dependences of κ_i on the atmospheric state variables $p(s)$, $T(s)$, and $x_1(s)$ to $x_N(s)$. We will discuss these dependences briefly here.

The most important dependence of κ_i is that on pressure $p(s)$. It comes from the fact that the width of the absorption line shape functions is governed by pressure broadening. An additional minor effect is that line positions also may depend on pressure, an effect known as pressure shift. To account for the pressure dependence of the κ_i , we have to store them on a pressure grid and interpolate if we need them for intermediate values. The interpolation is done in $\ln(p)$, because numerical experiments showed that this slightly improves the accuracy.

The second most important dependence of κ_i is that on temperature $T(s)$. Both the line widths and the line intensities depend on temperature. But only certain combinations of pressure and temperature occur in the Earth's atmosphere. Hence, storing the κ_i in a two dimensional table as a function of pressure and temperature would waste a lot of space. Instead, they are stored for a reference temperature profile $T^{\text{ref}}(p)$ and a set of temperature perturbations (in kelvin) for each pressure level. E.g., if the set of perturbations is $[-10, 0, +10]\text{K}$, then the κ_i are stored for three different temperatures for each pressure level: $[T^{\text{ref}}(p) - 10 \text{ K}, T^{\text{ref}}(p), T^{\text{ref}}(p) + 10 \text{ K}]$.

The least important dependence of κ_i is that on the trace gas VMRs $x_1(s)$ to $x_N(s)$. Some earlier absorption lookup table implementations have completely ignored this effect, as discussed in Section 2. However, in principle the width of an absorption line depends not only on total pressure, but also on the partial pressure of one or more trace gases. In theory this is always the case, because the broadening is different for each combination of collision partners. However, in practice trace gas concentrations in the Earth's atmosphere are normally so low, that this can be safely neglected. An important exception is water vapor in the lower troposphere, which can reach quite high volume mixing ratios. Therefore, the effect of water vapor mixing ratio on water vapor absorption (self

broadening), as well as on oxygen absorption (for example according to the parameterization by [24]) is not negligible.

To allow for this, the ARTS absorption lookup table can store κ_i also as a function of water vapor VMR ($x_{\text{H}_2\text{O}}$). This should not be done for all species i , but only for those where such a dependence is implemented in the line-by-line or continuum absorption model used. The strategy used to store the water vapor dependent κ_i is the same as for the temperature variations. I.e., κ_i are stored for a reference $x_{\text{H}_2\text{O}}$ profile and a set of perturbations.

In contrast to the temperature perturbations, fractional units are used for the $x_{\text{H}_2\text{O}}$ perturbations. Thus, a perturbation vector of $[0, 1, 10]$ means that κ_i is stored for a completely dry atmosphere, one with the reference $x_{\text{H}_2\text{O}}$ profile, and one with a 10 times increased $x_{\text{H}_2\text{O}}$ content.

All the κ_i , along with the pressure grid, the reference profiles of T and $x_{\text{H}_2\text{O}}$, and the perturbation vectors for T and $x_{\text{H}_2\text{O}}$ are stored in ARTS in a common structure, called absorption lookup table. There is also a frequency dimension in the table, with an associated frequency grid. No interpolation is done in that dimension.

3.2. Interpolation orders

In contrast to the absorption coefficients α , the κ_i are more closely related to the physics of absorption, rather than to the distribution of temperature and trace gases in the atmosphere. As a consequence, the dependence of the κ_i on p , T , and $x_{\text{H}_2\text{O}}$ can be described by smooth functions with few turning points. This means that it is advantageous to use higher order polynomial interpolation together with the lookup table, rather than simple linear interpolation.

To be completely clear here, N -th order interpolation means that a polynomial of degree N ($y = a_N x^N + a_{N-1} x^{N-1} + \dots + a_1 x + a_0$) is fitted to the data at the $N+1$ grid points closest to the interpolation point. First order interpolation corresponds to a plain linear interpolation between the two closest neighbours.

Appropriate interpolation orders were found by systematically comparing lookup table calculations with different interpolation orders to line-by-line calculations, and recording the achieved accuracies. The recommended interpolation orders resulting from this exercise are fifth order for p , seventh order for T , and fifth order for $x_{\text{H}_2\text{O}}$. These orders are set as defaults in ARTS, but can be changed by the user if desired. Generally, higher interpolation orders will lead to better accuracy, at the cost of somewhat higher computational cost of extraction. It is also possible to set the interpolation order to one minus the number of grid points in a given dimension. In that case the polynomial interpolation degenerates to a global polynomial fit in that dimension.

3.3. Interpolation algorithm

Efficient and general polynomial interpolation functions were developed and are now included in ARTS. The mathematics follows the treatment in [25], but the implementation differs to increase efficiency. Interpolation is done in three

steps: (1) Find the correct grid position in each dimension, (2) calculate weights, and (3) apply the weights to the data that should be interpolated. This implementation has the advantage of being very flexible, and that the weights can be re-used if the same interpolation has to be applied to different data elements. Note that the weight calculation step (2) also works for simultaneous polynomial interpolation in multiple dimensions, where the interpolation order in each dimension can be chosen independently.

3.4. Extraction algorithm

Absorption cross sections are extracted from the lookup table as a function of $\ln(p)$, T , and $x_{\text{H}_2\text{O}}$. A flowchart

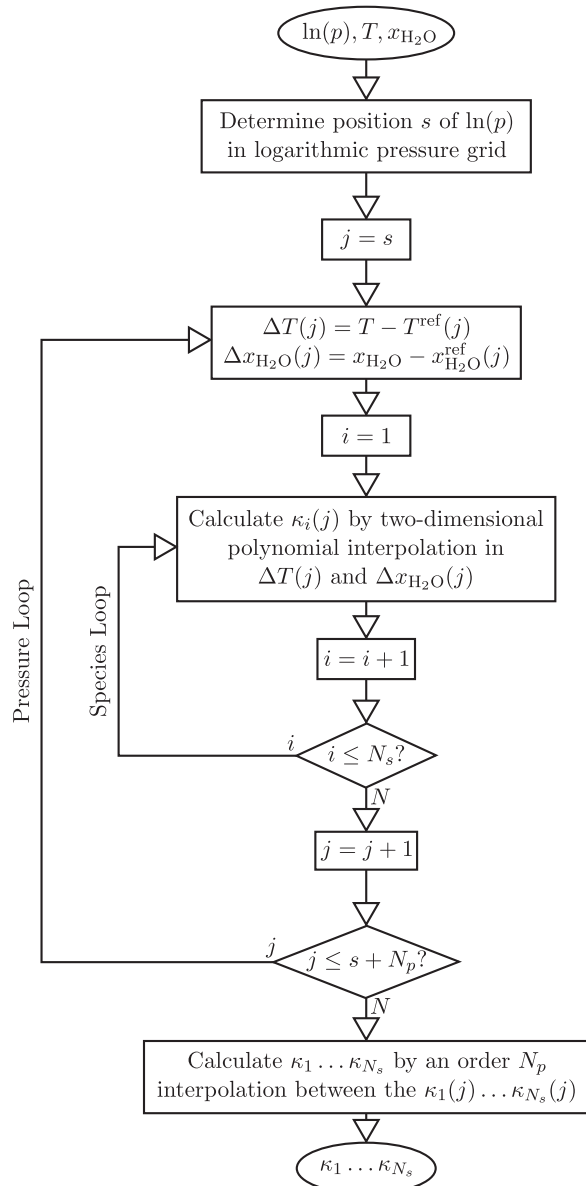


Fig. 1. A flowchart of the extraction algorithm for absorption cross-sections.

representation of the algorithm is shown in Fig. 1. The extraction algorithm is as follows:

1. Determine the position of $\ln(p)$ in the logarithmic pressure grid.
2. According to the interpolation order N_p in p , N_p+1 neighbours have to be used for the pressure interpolation. Do the following for all these $j = 1$ to N_p+1 pressure levels:
 - (a) Calculate $\Delta T(j) = T - T^{\text{ref}}(j)$.
 - (b) Calculate $\Delta x_{\text{H}_2\text{O}}(j) = x_{\text{H}_2\text{O}} - x_{\text{H}_2\text{O}}^{\text{ref}}(j)$.
 - (c) Calculate $\kappa_i(j)$ for all N_p+1 pressure levels by two-dimensional polynomial interpolation in $\Delta T(j)$ and $\Delta x_{\text{H}_2\text{O}}(j)$, with the appropriate interpolation orders N_T and $N_{x_{\text{H}_2\text{O}}}$.
3. Obtain κ_i by doing the order N_p interpolation between the $\kappa_i(j)$.

To get absorption coefficients α_i , the extracted κ_i only have to be multiplied by the number densities n_i , which are obtained with the ideal gas law from the x_i and p .

The description above is rather schematic, and omits details and special cases, such as the fact that $x_{\text{H}_2\text{O}}$ perturbations are only present for some species and not for others. Also, the extraction is typically done for all frequencies at the same time, giving as output a κ_i spectrum.

Obviously, the extraction is far from trivial to program (although not demanding in computation time). This is the price to pay for storing the κ_i in an irregular table, using deviations from reference profiles, instead of using simple Cartesian grids. However, that approach was judged necessary to make the lookup table efficient in terms of memory and computation time requirements.

3.5. Setting up lookup table generation

An important practical shortcoming in using any absorption lookup table is that it has to be known beforehand for which range of p , T , and $x_{\text{H}_2\text{O}}$ absorption coefficients will be needed. To relieve the user from the burden of having to define these ranges explicitly, ARTS includes automatic setup routines for the lookup table for three important cases: '3D' a 3D atmosphere, 'batch' a batch of (1D or 3D) calculations, and 'wide' a wide table setup that should be suitable for all reasonable atmospheric conditions.

In the first two cases, the setup routine analyzes the actual atmospheric states for which the simulations will be carried out. It calculates reference profiles of p , T , and $x_{\text{H}_2\text{O}}$ as simple mean profiles. (The mean is either across the 3D field, or across the calculation batch.) The setup routine then determines the minimum and maximum profiles of temperature and water vapor, and uses them to construct temperature and water vapor perturbation vectors with a user-defined or default spacing. The set of trace gases, for which $x_{\text{H}_2\text{O}}$ perturbations should be taken into account, is also set automatically, based on a fixed internal list. Currently this list contains only water vapor itself and oxygen, the latter only if it includes a continuum model, which may depend on water vapor concentration.

The default p grid spacing is 0.05 in $\log(\text{Pa})$ units. The default T grid spacing is 20 K, and the default $x_{\text{H}_2\text{O}}$ grid spacing 100 in fractional units. If these default spacings result in fewer grid points than required by the interpolation order, they are refined, so that each grid has at least $N+1$ points for interpolation order N . This means in practice that for $x_{\text{H}_2\text{O}}$ the actual grid spacing will be determined purely by the interpolation order, since the default value is large and will always require refinement. These default settings result in very accurate calculations, as will be shown in the next section.

The last setup routine ‘wide’ is intended for cases where the atmospheric states are not known beforehand. There we use constant reference profiles of T and $x_{\text{H}_2\text{O}}$, so that the different table dimensions really are orthogonal, not skewed as in the general case. The default values for this case are chosen such that a wide range of reasonable atmospheric conditions is covered (pressure between 0.5 Pa and 1100 hPa, temperature between 100 and 400 K, and humidity between 0 and 5% VMR). The default grid spacings are similar to the ‘batch’ case. The disadvantage of this last setup option is that the table will typically be roughly twice as large as in the ‘batch’ case, since a wider range of atmospheric conditions is covered.

As a consequence of these automatic setup methods, it is very simple to switch ARTS between generating (and using) the lookup table on one hand, and line-by-line calculation ‘on the fly’ on the other hand. This flexibility has been used extensively for lookup table validation, as described in the next section.

4. Validation

We validated the lookup table method by comparing radiative transfer simulations with absorption lookup table to reference simulations, for which absorption was calculated on the fly by a line-by-line calculation. In subsequent sections, we will sometimes use ‘(tab)’ to refer to calculations with lookup table, and ‘(lbl)’ to refer to calculations with on the fly line-by-line calculation.

The ‘on the fly’ option means that absorption is calculated line-by-line from the local pressure, temperature, and trace gas VMRs. This is done during the integration of the radiative transfer equation whenever absorption for a point in the atmosphere is needed. Thus, with this option there is no interpolation at all in the absorption coefficients. Comparing simulated radiances between lookup table and on the fly line-by-line simulations is thus a true check of all errors associated with the use of the lookup table.

This comparison was done for a large set of atmospheric states and for different instrument scenarios. Details are described in the following sections.

4.1. Atmospheric states

We used 1000 different atmospheric states for testing, which were randomly selected from the q , T , and O_3 datasets of [26], the same 1000 cases that were used in [27]. For each state, the Chevallier dataset contains atmospheric profiles on 91 vertical levels. For our analysis, we

used the profiles of pressure, altitude, temperature, water vapor concentration, and ozone concentration. Oxygen and nitrogen concentrations were assumed to be fixed, with VMRs of 0.2095 and 0.7808, respectively.

Hydrometeors (cloud liquid water and cloud ice water) were ignored, only clear-sky radiative transfer simulations were done. The reason for this is that calculation with hydrometeor scattering with the ‘on the fly’ option are prohibitively expensive. (Originally, this was the main reason for developing the lookup table approach.)

4.2. Instrument scenarios

Three different instrument scenarios were used: AMSU-B, HIRS, and Odin. These represent to some extent the range of instruments for which ARTS is typically used. In all three cases, the tests used the standard configuration files for these instruments, which are part of the ARTS distribution. Also, the lookup table setup in all cases was simply using the ARTS defaults, identical to the settings described in Section 3.

The first tested sensor, AMSU-B, is a down-looking millimeter-wave satellite sensor with two window channels and three channels centered on the 183 GHz water vapor absorption line. By convention, channels are denoted as Channel 16–20, leaving the lower numbers for AMSU-A. The sensor is described in more detail for example in [28] and in even more technical detail in [29]. All five channels were tested. As viewing geometry we arbitrarily chose nadir, but the results for other looking angles are expected to be very similar, as far as the accuracy of the lookup table is concerned (not the actual simulated radiances). The standard ARTS setup for AMSU-B includes ozone as an absorbing species, since it was shown by [30] that ozone has a non-negligible impact in the innermost water vapor channel (Channel 18).

The second tested sensor, HIRS, is a down-looking infrared satellite sensor with 12 channels in the thermal infrared spectral range and seven channels in the near infrared spectral range. The sensor is described in more detail for example in [27], or, in its original version, in [31]. We used only the thermal infrared channels for the test, denoted as Channels 1–12. The instrument setup in ARTS, which was used for the test calculations, uses the method of representative frequencies to efficiently simulate each HIRS channel with few monochromatic calculations. The frequencies were derived with a simulated annealing algorithm in [27]. As in the AMSU-B case, the test was done only for the nadir viewing direction, as the lookup table performance should not depend significantly on the instrument viewing angle.

The third tested sensor, Odin, is a millimeter/sub-millimeter limb sounder. Some rough information on the sensor is given in [10], more detailed information is given in [32]. For the test, limb spectra were simulated for some different tangent altitudes between 20 and 80 km. The Odin bands at 119 and 501 GHz were used. In the latter case, minor trace gases were ignored in the simulation, although they have signatures in the band, because no input profiles were readily available. This has no impact on the validity of the test, since water vapor,

oxygen, and ozone generate enough spectral structure to make the test representative of general limb sounder simulations.

4.3. Test results and discussion

We simulated each instrument scenario for each of the 1000 atmospheric cases with and without lookup table. Then we looked at the statistics of $d = T_b(\text{tab}) - T_b(\text{lbl})$, the difference in the radiance (in brightness temperature units) between the two calculation alternatives. In particular, we looked at the mean value \bar{d} and standard deviation σ_d over the 1000 cases. These are summarized in Table 2. In the case of Odin, which is a spectrally resolving instrument, the averaging was done also over frequency (inside each band), and over the different tangent altitudes.

The two setup options ‘batch’ and ‘wide’ (see Section 3) were tested separately. The table shows that both options produce very accurate results. Note that the unit for brightness temperature differences in the table is microkelvin, so the largest standard deviation found (for the Odin 501 GHz band) is only 0.006 K, and the largest mean error (for the same satellite and band) is only 0.009 K.

Basically, the table shows that the errors associated with the use of the lookup table are negligible compared to instrument noise, which is typically of the order of 1 K. Other forward model errors, for example due to imperfectly known spectroscopic parameters or continua, or

even due to different interpolation strategies in the radiative transfer, are typically of the same 1 K order [33].

Interestingly, the Odin 501 GHz band, in connection with the ‘wide’ setup option, produces by far the largest errors. (All other errors are well below 1 mK.) The reason for this is the constant $x_{\text{H}_2\text{O}}$ reference profile with the ‘wide’ option, which is not optimal for the stratosphere, where water vapor concentrations are very low. This effect is not noticeable for the 119 GHz band, which is dominated by oxygen absorption, but clearly noticeable for the 501 GHz band. It would not be difficult to optimize this case further, but the error is still so small that we felt no need to do this.

Total calculation times for the 1000 atmospheric cases are listed in Table 3. Both the actual time that has passed (wallclock time) and the CPU time are given. CPU times are significantly higher, since the jobs were run on eight-core machines, and ARTS uses Open MP parallelization to speed up the calculations.

The last column in the table shows the speedup factor, defined as run time with ‘on the fly’ absorption ($\Delta t(\text{lbl})$) divided by run time with absorption lookup table ($\Delta t(\text{tab})$). As expected, the speedup factor depends strongly on the cost of the line-by-line calculation. It is largest for HIRS, where more than a million lines are considered, and smallest for Odin, where only a few lines are considered. The table shows also that the speedup factor in wallclock time is less than in CPU time. This is because wallclock time depends on external factors, not the least computer load, which are not influenced by the lookup table.

The numbers given here are for the ‘batch’ setup case, numbers for the ‘wide’ setup are qualitatively similar. The numbers given for $\Delta t(\text{tab})$ include the time for the lookup table generation. Because of this, speedup factors will also depend on the number of atmospheric cases. For situations where only very few atmospheric cases have to be calculated, it may not be worth to pre-calculate the lookup table, and the ‘on the fly’ option may actually be computationally cheaper. We make no attempt here to calculate the break even point, i.e., the number of cases that should be exceeded to make the lookup table more efficient than ‘on the fly’. It depends on too many factors, not only on the cost of the line-by-line calculation, but also on details of the radiative transfer calculation that determine how often an absorption calculation is needed.

Table 2

Lookup table test results, based on 1000 atmospheric cases. For each case, $d = T_b(\text{tab}) - T_b(\text{lbl})$ was calculated as the difference between the simulated measurement with lookup table and with explicit line-by-line calculation. Column ‘ \bar{d} ’ states the mean of d , column ‘ σ_d ’ its standard deviation. All numbers are for brightness temperature in units of microkelvin, so the differences in all cases are small compared to the measurement noise. In the case of Odin, which is a spectrally resolving instrument, the average and standard deviation are not only over atmospheric case, but also over tangent altitude and frequency within the band. ‘Batch’ and ‘Wide’ are the two different default table setup methods included in ARTS.

Instrument	Channel	‘Batch’ setup		‘Wide’ setup	
		\bar{d} (μK)	σ_d (μK)	\bar{d} (μK)	σ_d (μK)
AMSU	16	374.4	855.7	-177.4	174.5
	17	100.1	546.4	-165.2	140.2
	18	12.9	68.0	0.5	1.7
	19	-160.1	311.1	-0.3	2.7
	20	193.4	291.8	-14.9	30.8
HIRS	1	-0.6	1.2	-447.3	220.1
	2	0.1	0.3	-250.9	316.4
	3	-0.1	0.3	8.9	199.9
	4	0.1	3.2	494.9	185.2
	5	-2.2	4.7	758.5	207.5
	6	-3.7	9.4	884.2	274.8
	7	-18.7	21.5	267.9	105.1
	8	-17.9	37.9	6.3	8.8
	9	-8.7	17.6	14.2	8.1
	10	-10.1	25.8	11.8	10.0
	11	-6.1	7.7	3.9	0.8
	12	-2.1	4.4	-5.8	4.6
Odin	119GHz	-7.5	53.7	-37.7	68.4
	501GHz	-23.5	49.3	-8847.0	5823.7

Table 3

ARTS program run times with lookup table and ‘on the fly’ absorption. 1000 atmospheric profiles were processed. The $\Delta t(\text{tab})$ times include the time for the lookup table generation. The last column gives the speedup factor, see text for definition.

Instrument	Time type	$\Delta t(\text{lbl})$ (s)	$\Delta t(\text{tab})$ (s)	Factor
AMSU	CPU	22,331	849	26
	Wallclock	2831	125	23
HIRS	CPU	382,938	1,662	230
	Wallclock	48,990	326	150
Odin (501 GHz)	CPU	11,182	6231	1.8
	Wallclock	1447	924	1.6

The interested user is recommended to simply try both options and see which is faster.

Since the errors with the default setup settings are very small, the user could in principle modify these settings to reduce memory consumption or to further increase calculation speed, if a poorer accuracy is acceptable. However, this is not completely straightforward, since there are in total six parameters to consider: the grid spacing in p , T , and $x_{\text{H}_2\text{O}}$, and the three associated interpolation orders. We thus recommend to use the default settings.

5. Summary and conclusions

This article describes the lookup table approach that is used to store pre-calculated absorption data in the radiative transfer model ARTS. The table stores absorption cross-sections as a function of frequency, pressure, temperature, and the water vapor volume mixing ratio, where the last dimension is only included for those gas species that require it. The table is used together with an extraction strategy, which uses polynomial interpolation, with recommended interpolation orders between five and seven for p , T , and $x_{\text{H}_2\text{O}}$. No interpolation is done in frequency. We also derived recommended default settings for grid spacings and interpolation orders, and verified that the approach gives very accurate results with these default settings. Errors introduced by the lookup table were always below a few millikelvin, in terms of the simulated brightness temperature.

The main advantage of using the lookup table is the significant calculation speed increase, with speedup factors exceeding 200 for infrared calculations with many spectral lines. (The exact speedup factor depends on the number of atmospheric cases calculated, if one includes the one-time computation time to generate the lookup table.) Are there also disadvantages? Yes, two from a user perspective. Firstly, one has to anticipate to some degree the range of atmospheric conditions for which calculations have to be carried out, since the lookup table has to be pre-calculated. We try to help with the table setup by providing the automatic setup methods that are discussed in Section 3.5. This is easy for 'well behaved' input data, but can be practically difficult for irregular data, such as radiosonde data, which can have unphysical values, such as negative temperatures, or odd jumps in pressure. Thus, with lookup table, ARTS is somewhat less robust against such outliers or odd values in input data.

The second disadvantage is that the lookup table consumes quite a lot of memory (typically some tens or even hundreds of megabytes). We do not normally recommend to store the tables, but to generate them at the beginning of each batch run. For a table intended for permanent storage, and to be passed on to other people, it would be better to use an internal compression method, such as the one developed by [19]. We have so far not encountered any practical obstacles due to the large table size, even the calculation of total outgoing longwave radiation fluxes or high-resolution reference calculations for HIRS with many 1000 frequencies per channel are feasible with current PCs. We have therefore not implemented any compression so far. But it would

certainly be possible, following the approach by Strow and coworkers.

All in all, the calculation speed increase will outweigh the two mentioned disadvantages for most users. To use the absorption lookup table is therefore now the recommended default option in ARTS, rather than 'on the fly' absorption calculation.

Acknowledgements

We acknowledge the help of the ARTS radiative transfer community, many of whom have indirectly contributed by implementing features to the ARTS model.

Last but not least, we acknowledge our funding agencies. The ESA study 'Development of a Radiative Transfer Model for Frequencies between 200 and 1000 GHz' (Contract 17632/03/NL/FF) was especially valuable. Furthermore, direct or indirect support has been provided by the Swedish National Space Board (e.g., Contracts 59/07, 92/08, 48/09) and by the Swedish Research Council (Contracts 2007-3720 and 2007-5370).

References

- [1] Buehler SA, Eriksson P, Kuhn T, von Engeln A, Verdes C. ARTS, the atmospheric radiative transfer simulator. *J Quant Spectrosc Radiat Transfer* 2005;91(1):65–93, doi:10.1016/j.jqsrt.2004.05.051.
- [2] Eriksson P, Buehler SA, Davis CP, Emde C, Lemke O. ARTS, the atmospheric radiative transfer simulator, version 2. *J Quant Spectrosc Radiat Transfer* 2011; this issue, doi:10.1016/j.jqsrt.2011.03.001.
- [3] Eriksson P, Jiménez C, Buehler SA. Qpack, a general tool for instrument simulation and retrieval work. *J Quant Spectrosc Radiat Transfer* 2005;91(1):47–64, doi:10.1016/j.jqsrt.2004.05.050.
- [4] Buehler SA, von Engeln A, Brocard E, John VO, Kuhn T, Eriksson P. Recent developments in the line-by-line modeling of outgoing longwave radiation. *J Quant Spectrosc Radiat Transfer* 2006;98(3):446–57, doi:10.1016/j.jqsrt.2005.11.001.
- [5] John VO, Buehler SA, von Engeln A, Eriksson P, Kuhn T, Brocard E, et al. Understanding the variability of clear-sky outgoing long-wave radiation based on ship-based temperature and water vapor measurements. *Q J R Meteorol Soc* 2006;132(621):2675–91, doi:10.1256/qj.05.70.
- [6] Saunders R, Matricardi M, Brunel P. An improved fast radiative transfer model for assimilation of satellite radiance observations. *Q J R Meteorol Soc* 1999;125:1407–25.
- [7] Matricardi M. Technical note: an assessment of the accuracy of the RTTOV fast radiative transfer model using IASI data. *Atmos Chem Phys* 2009;9(2):9491–535, doi:10.5194/acp-9-6899-2009.
- [8] Buehler SA, Courcoux N, John VO. Radiative transfer calculations for a passive microwave satellite sensor: Comparing a fast model and a line-by-line model. *J Geophys Res* 2006;111:D20304, doi:10.1029/2005JD006552.
- [9] Buehler SA, Jiménez C, Evans KF, Eriksson P, Rydberg B, Heymsfield AJ, et al. A concept for a satellite mission to measure cloud ice water path and ice particle size. *Q J R Meteorol Soc* 2007;133(S2):109–28, doi:10.1002/qj.143.
- [10] Rydberg B, Eriksson P, Buehler SA, Murtagh DP. Non-Gaussian Bayesian retrieval of tropical upper tropospheric cloud ice and water vapour from Odin-SMR measurements. *Atmos Meas Tech* 2009;2:621–37, doi:10.5194/amt-2-621-2009.
- [11] Davis C, Emde C, Harwood R. A 3D polarized reversed Monte Carlo radiative transfer model for mm and sub-mm passive remote sensing in cloudy atmospheres. *IEEE Trans Geosci Remote* 2005;43(5):1096–101, doi:10.1109/TGRS.2004.837505.
- [12] Davis CP, Evans KF, Buehler SA, Wu DL, Pumphrey HC. 3-D polarised simulations of space-borne passive mm/sub-mm mid-latitude cirrus observations: a case study. *Atmos Chem Phys* 2007;7:4149–58, doi:10.5194/acp-7-4149-2007.
- [13] Emde C, Buehler SA, Davis C, Eriksson P, Sreerekha TR, Teichmann C. A polarized discrete ordinate scattering model for simulations of limb

- and nadir longwave measurements in 1D/3D spherical atmospheres. *J Geophys Res* 2004;109(D24):D24207, doi:[10.1029/2004JD005140](https://doi.org/10.1029/2004JD005140).
- [14] Emde C, Buehler SA, Eriksson P, Sreerekha TR. The effect of cirrus clouds on microwave limb radiances. *J Atmos Res* 2004;72(1–4): 383–401, doi:[10.1016/j.atmosres.2004.03.023](https://doi.org/10.1016/j.atmosres.2004.03.023).
- [15] Koukoulis ME, Irwin PGJ, Taylor FW. Water vapor abundance in Venus' middle atmosphere from Pioneer Venus OIR and Venera 15 FTS measurements. *Icarus* 2005;173(1):84–99, doi:[10.1016/j.icarus.2004.08.023](https://doi.org/10.1016/j.icarus.2004.08.023).
- [16] Scott NA, Chedin A. A fast line-by-line method for atmospheric absorption computations: the automatized atmospheric absorption Atlas. *J Appl Meteorol* 1981;20(7):802–12.
- [17] Aoki T. Development of a line-by-line model for the infrared radiative transfer in the Earth's atmosphere. *Pap Meteorol Geophys* 1988;39(2):53–8.
- [18] Turner DS. Absorption coefficient estimation using a two-dimensional interpolation procedure. *J Quant Spectrosc Radiat Transfer* 1995;53(6):633–7.
- [19] Strow LL, Motteler HE, Benson RG, Hannon SE, Souza-Machado SD. Fast computation of monochromatic infrared atmospheric transmittances using compressed look-up tables. *J Quant Spectrosc Radiat Transfer* 1998;59(3–5):481–93, doi:[10.1016/S0022-4073\(97\)00169-6](https://doi.org/10.1016/S0022-4073(97)00169-6).
- [20] Mitsel AA, Tashkun SA, Okladnikov IG, Milyakov AV. Methodological problems of compiling the data bank of atmospheric gas absorption coefficients. *Proc SPIE* 2001;4341:616–25 doi:[10.1117/12.412007](https://doi.org/10.1117/12.412007). Seventh International Symposium on Atmospheric and Ocean Optics; Tomsk; 19 July 2000 through 22 July 2000; Code 58035.
- [21] Dudhia A, Morris PE, Wells RJ. Fast monochromatic radiative transfer calculations for limb sounding. *J Quant Spectrosc Radiat Transfer* 2002;74(6):745–56, doi:[10.1016/S0022-4073\(01\)00285-0](https://doi.org/10.1016/S0022-4073(01)00285-0).
- [22] Clough SA, Shephard MW, Worden J, Brown PD, Worden HM, Luo M, et al. Forward model and Jacobians for tropospheric emission spectrometer retrievals. *IEEE Trans Geosci Remote* 2006;44(5): 1308–23, doi:[10.1109/TGRS.2005.860986](https://doi.org/10.1109/TGRS.2005.860986).
- [23] Goody RM, Yung YL. Atmospheric radiation theoretical basis. 2nd ed.. Oxford University Press; 1995 ISBN 0-19-505134-3.
- [24] Rosenkranz PW. Absorption of microwaves by atmospheric gases. In: Janssen MA, editor. Atmospheric remote sensing by microwave radiometry. John Wiley and Sons, Inc.; 1993. p. 37–90. ISBN: 0-471-62891-3, <ftp://mesa.mit.edu/phil/lbl_rt>.
- [25] Press WH, Teukolsky SA, Vetterling WT, Flannery BP. Numerical recipes in C. 2nd ed.. Cambridge University Press; 1992 ISBN 0-521-43108-5.
- [26] Chevallier F, Di Michele S, McNally AP. Diverse profile datasets from the ECMWF 91-level short-range forecasts. Technical Report; NWP SAF Satellite Application Facility for Numerical Weather Prediction; 2006. Document No. NWPSAF-EC-TR-010, Version 1.0.
- [27] Buehler SA, John VO, Kottayil A, Milz M, Eriksson P. Efficient radiative transfer simulations for a broadband infrared radiometer—combining a weighted mean of representative frequencies approach with frequency selection by simulated annealing. *J Quant Spectrosc Radiat Transfer* 2010;111(4):602–15, doi:[10.1016/j.jqsrt.2009.10.018](https://doi.org/10.1016/j.jqsrt.2009.10.018).
- [28] Buehler SA, Kuvatov M, John VO, Milz M, Soden BJ, Jackson DL, et al. An upper tropospheric humidity data set from operational satellite microwave data. *J Geophys Res* 2008;113:D14110, doi:[10.1029/2007JD009314](https://doi.org/10.1029/2007JD009314).
- [29] Saunders RW, Hewison TJ, Stringer SJ, Atkinson NC. The radiometric characterization of AMSU-B. *IEEE Trans Microw Theory* 1995;43(4): 760–71.
- [30] John VO, Buehler SA. The impact of ozone lines on AMSU-B radiances. *Geophys Res Lett* 2004;31:L21108, doi:[10.1029/2004GL021214](https://doi.org/10.1029/2004GL021214).
- [31] Smith WL, Woolf HM, Hayden CM, Wark DQ, McMillin LM. The TIROS-N operational vertical sounder. *Bull Am Meteorol Soc* 1979;60:1177–87.
- [32] Murtagh D, Frisk U, Merino F, Ridal M, Jonsson A, Stegman J, et al. An overview of the Odin atmospheric mission. *Can J Phys* 2002;80(4):309–19.
- [33] Melsheimer C, Verdes C, Buehler SA, Emde C, Eriksson P, Feist DG, et al. Intercomparison of general purpose clear sky atmospheric radiative transfer models for the millimeter/submillimeter spectral range. *Radio Sci* 2005; RS1007 doi:[10.1029/2004RS003110](https://doi.org/10.1029/2004RS003110).