CKDMIP: Technical Guide

Robin J. Hogan
European Centre for Medium Range Weather Forecasts, Reading, UK

Document version 0.7 (28 October 2019) applicable to CKDMIP software version 0.7*
CKDMIP home page: https://confluence.ecmwf.int/display/CKDMIP

1 Introduction

The motivation and overall structure of the Correlated K-Distribution Model Intercomparison Project (CKDMIP) is described in detail by Hogan and Matricardi (2020). This document describes the CKDMIP datasets (section 2), software (section 3) and what files should be provided by project participants (section 4).

2 Datasets

The CKDMIP datasets are available at `ftp://dissemination.ecmwf.int`, username `ckdmip`, password available on request from Robin Hogan (`r.j.hogan@ecmwf.int`). Some of the smaller files are also available for direct download from the CKDMIP home page.

All files are in NetCDF format. The smaller files use the NetCDF3 'classic' format (suffix `nc`), which is compatible with a wider range of software. The larger files use NetCDF4 format, which is actually HDF5 on disk (suffix `h5`), but can be accessed using the same NetCDF function calls provided that your software is compiled against version 4 of the NetCDF library. The latter format is required because of its support for very large files and per-variable compression.

Table 3 of Hogan and Matricardi (2020) describes the four dataset groups provided in CKDMIP. The first two, 'Evaluation-1' and 'Evaluation-2' each contain 50 realistic atmospheric profiles, with the first intended to be used for training and evaluation of CKD models, and the second for independent evaluation. At the time of writing, 'Evaluation-2' has not been released. A further two, 'MMM' and 'Idealized', provide additional profiles that may be useful to some participants for generating CKD models.

2.1 Gas concentrations

The following files in the directory `concentrations` on the FTP site:

```
ckdmip_evaluation1_concentrations.nc
ckdmip_mmm_concentrations.nc
ckdmip_idealized_concentrations.nc
```

contain the mole fractions (in units of mol mol^{-1}, equivalent to volume mixing ratio in m^3 m^{-3}) for the nine gases considered in CKDMIP: H_2O, O_3, N_2, O_2, CO_2, CH_4, N_2O, CFC11 and CFC12. They are provided on both on 'full levels' (i.e. layer-averaged values) and 'half levels' (the interfaces between layers). If a gas is 'well mixed', either having a constant mole fraction with height or varying as a prescribed function of pressure, then the mole fraction variables also have a `reference_surface_value` attribute containing the nominal surface value. This

then enables the software described in section 3 to scale the spectral optical depths to correspond to a different concentration. The files also contain the temperature and pressure on full and half levels.

Note that half-level quantities are the primary variables: these are the ones used by the Line-By-Line Radiative Transfer Model (LBLRTM; Clough et al, 2005), and the full-level quantities were derived from the half-level quantities: for pressure by taking the average of the two enclosing half-level values, and for temperature and mole fraction by performing a pressure-weighted average assuming a linear variation with pressure within the layer.

As described by Hogan and Matricardi (2020), the 'Evaluation-1' dataset contains 50 realistic profiles (the `column` dimension is 50) of temperature, H_2O and O_3 extracted from the ECMWF model, along with present-day (i.e. the year 2020) profiles for the other gases that are either constant with pressure (for N_2 and O_2) or vary as a prescribed function of pressure.

The `column` dimension of the 'MMM' dataset is 3, with these profiles containing the median, minimum and maximum temperatures from the original ECMWF dataset. For seven of the nine gases, a single concentration profile is provided for each column (on both full and half levels), which is in fact constant with column and corresponds to present-day conditions. But for H_2O and O_3, six variables are provided for each gas: for H_2O we have `h2o_median_mole_fraction_hl`, `h2o_minimum_mole_fraction_hl` and `h2o_maximum_mole_fraction_hl`, corresponding to the median, minimum and maximum concentrations of the original ECMWF dataset, on half levels. The 'maximum' concentrations have been capped at liquid water saturation. Each has values for the three columns (corresponding to varying temperature). The corresponding three full-level variables have the suffix `fl`, and similarly for O_3.

The 'Idealized' dataset contains 6 columns corresponding to temperatures 20-K apart. All of the mole-fraction variables are constant with pressure and temperature. Since the molar absorption of water vapour is dependent on temperature, for this gas alone 12 separate concentration profiles are provided (named `h2o_a_mole_fraction_hl` to `h2o_l_mole_fraction_hl`).

2.2 Absorption spectra

For each mole-fraction variable in the concentration files described in section 2.1, LBLRTM has been used to compute a profile of the spectral optical depth for that one gas. To keep the size of individual files less than around 10 GB, the 50 profiles of the 'Evaluation-1' dataset (in the `lw_spectra/evaluation1` directory on the FTP site) have been split into five groups of 10, with the first group having file names as follows:

```
ckdmip_evaluation1_lw_spectra_h2o_present_1-10.h5
ckdmip_evaluation1_lw_spectra_o3_present_1-10.h5
ckdmip_evaluation1_lw_spectra_n2_constant_1-10.h5
ckdmip_evaluation1_lw_spectra_o2_constant_1-10.h5
ckdmip_evaluation1_lw_spectra_co2_present_1-10.h5
ckdmip_evaluation1_lw_spectra_ch4_present_1-10.h5
ckdmip_evaluation1_lw_spectra_n2o_present_1-10.h5
ckdmip_evaluation1_lw_spectra_cfc11_present-equivalent_1-10.h5
ckdmip_evaluation1_lw_spectra_cfc12_present_1-10.h5
```

All files 45 amount to around 228 GB of data in total. NetCDF4/HDF5 compression has been used, with the result that the gases with the most spectral features have the largest file sizes. In the file names, `constant` indicates a profile whose mole fraction is constant with pressure, `present` indicates a present-day profile, and `present-equivalent` indicates that the concentration of CFC11 has been increased to approximately represent the radiative forcing of 38 more minor greenhouse gases (Hogan and Matricardi, 2020).

Each file contains the pressure, temperature and mole fraction on full and half levels, for the 10 profiles, which have basically been copied from the corresponding concentration file described in section 2.1. The `wavenumber` variable defines 7211999 wavenumbers in units of cm^{-1}, with a spacing of 0.0002 cm^{-1} in the spectral range 0–1300 cm^{-1}, 0.001 cm^{-1} in the spectral range 1300–1700 cm^{-1}, and 0.005 cm^{-1} in the spectral range 1700-3260 cm^{-1}.

The `optical_depth` variable then provides the optical depth of each atmospheric layer at each wavenumber. If needed, the molar absorption coefficient (in $\mathrm{m}^2\ \mathrm{mol}^{-1}$) of layer i, k_i, could be computed from the layer optical

depth, δ_i, using

$$k_i = \frac{g M_d \delta_i}{x_i \left(p_{i+1/2} - p_{i-1/2} \right)}, \tag{1}$$

where g is the acceleration due to gravity, $M_d = 0.02897$ kg mol^{-1} is the molar mass of dry air, x_i is the layer-mean mole fraction of the gas in question in layer i, and $p_{i+1/2}$ is the pressure of the half level between layers i and $i+1$. All quantities have SI units.

The directory `lw_spectra/mmm` contains the absorption spectra for the 'MMM' dataset, with separate files for the three different concentration scenarios of H_2O and O_3 (total around 24 GB):

```
ckdmip_mmm_lw_spectra_h2o_median.h5
ckdmip_mmm_lw_spectra_h2o_minimum.h5
ckdmip_mmm_lw_spectra_h2o_maximum.h5
ckdmip_mmm_lw_spectra_o3_median.h5
ckdmip_mmm_lw_spectra_o3_minimum.h5
ckdmip_mmm_lw_spectra_o3_maximum.h5
ckdmip_mmm_lw_spectra_n2_constant.h5
ckdmip_mmm_lw_spectra_o2_constant.h5
ckdmip_mmm_lw_spectra_co2_present.h5
ckdmip_mmm_lw_spectra_ch4_present.h5
ckdmip_mmm_lw_spectra_n2o_present.h5
ckdmip_mmm_lw_spectra_cfc11_present-equivalent.h5
ckdmip_mmm_lw_spectra_cfc12_present.h5
```

The directory `lw_spectra/idealized` contains the absorption spectra for the 'Idealized' dataset, with separate files for the 12 different concentrations of H_2O (total around 80 GB):

```
ckdmip_idealized_lw_spectra_h2o_constant-a.h5
ckdmip_idealized_lw_spectra_h2o_constant-b.h5
ckdmip_idealized_lw_spectra_h2o_constant-c.h5
ckdmip_idealized_lw_spectra_h2o_constant-d.h5
ckdmip_idealized_lw_spectra_h2o_constant-e.h5
ckdmip_idealized_lw_spectra_h2o_constant-f.h5
ckdmip_idealized_lw_spectra_h2o_constant-g.h5
ckdmip_idealized_lw_spectra_h2o_constant-h.h5
ckdmip_idealized_lw_spectra_h2o_constant-i.h5
ckdmip_idealized_lw_spectra_h2o_constant-j.h5
ckdmip_idealized_lw_spectra_h2o_constant-k.h5
ckdmip_idealized_lw_spectra_h2o_constant-l.h5
ckdmip_idealized_lw_spectra_o3_constant.h5
ckdmip_idealized_lw_spectra_n2_constant.h5
ckdmip_idealized_lw_spectra_o2_constant.h5
ckdmip_idealized_lw_spectra_co2_constant.h5
ckdmip_idealized_lw_spectra_ch4_constant.h5
ckdmip_idealized_lw_spectra_n2o_constant.h5
ckdmip_idealized_lw_spectra_cfc11_constant-equivalent.h5
ckdmip_idealized_lw_spectra_cfc12_constant.h5
```

2.3 Reference broadband fluxes

Section 3 describes how the CKDMIP software can be used to perform line-by-line radiative transfer calculations on the absorption spectra described in section 2.2 to generate reference fluxes (both spectral and broadband) against which to test CKD models. Since the code is slow to run, `lw_fluxes/evaluation1` of the FTP contains the broadband upwelling and downwelling fluxes for each of the relevant scenarios applied to the 50 profiles of the Evaluation-1 dataset, in the following files (total around 4 MB):

```
ckdmip_evaluation1_lw_fluxes_glacialmax.h5
ckdmip_evaluation1_lw_fluxes_preindustrial.h5
ckdmip_evaluation1_lw_fluxes_present.h5
ckdmip_evaluation1_lw_fluxes_future.h5
ckdmip_evaluation1_lw_fluxes_co2-140.h5
ckdmip_evaluation1_lw_fluxes_co2-280.h5
ckdmip_evaluation1_lw_fluxes_co2-560.h5
ckdmip_evaluation1_lw_fluxes_co2-1120.h5
ckdmip_evaluation1_lw_fluxes_co2-2240.h5
ckdmip_evaluation1_lw_fluxes_ch4-350.h5
ckdmip_evaluation1_lw_fluxes_ch4-700.h5
ckdmip_evaluation1_lw_fluxes_ch4-1200.h5
ckdmip_evaluation1_lw_fluxes_ch4-2600.h5
ckdmip_evaluation1_lw_fluxes_ch4-3500.h5
ckdmip_evaluation1_lw_fluxes_n2o-190.h5
ckdmip_evaluation1_lw_fluxes_n2o-270.h5
ckdmip_evaluation1_lw_fluxes_n2o-405.h5
ckdmip_evaluation1_lw_fluxes_n2o-540.h5
ckdmip_evaluation1_lw_fluxes_cfc11-0.h5
ckdmip_evaluation1_lw_fluxes_cfc11-218.h5
ckdmip_evaluation1_lw_fluxes_cfc11-2000.h5
ckdmip_evaluation1_lw_fluxes_cfc12-0.h5
ckdmip_evaluation1_lw_fluxes_cfc12-200.h5
ckdmip_evaluation1_lw_fluxes_cfc12-550.h5
```

As indicated in the file name, this includes the four scenarios (glacial maximum, preindustrial, present-day and future) listed in Table 2 of Hogan and Matricardi (2020). It also includes perturbations of individual gases keeping the others at present-day conditions, where the numbers in the filenames indicate ppmv for CO_2, ppbv for CH_4 and N_2O, and pptv for CFC11 and CFC12. These files also contain fluxes (in W m^{-2}) in the 13 'narrow' bands defined in Table 4 of Hogan and Matricardi (2020).

3 Software

The latest version of the software package, named `ckdmip-x.y.tar.gz`, may be obtained from the CKDMIP home page, and contains the `ckdmip_tool`, `ckdmip_lw` and `ckdmip_sw` programs. They are coded in Fortran and it have been tested under Linux. Instructions on compiling are provided in the `README` file within. The `ckdmip_tool` program performs various tasks described in section 3.1, such as creating Rayleigh-scattering and cloud spectral files, and creating solar spectral irradiance files. The `ckdmip_lw` program can be used in three ways: to produce reference line-by-line fluxes (section 3.2), to merge and store the absorption spectra of individual gases (section 3.3), or to produce fluxes for CKD models (section 3.4).

3.1 Utility for creating input files

The `ckdmip_tool` program performs various tasks according to the command-line arguments; typically it generates a file at full spectral resolution from information at coarser resolution. The arguments available to all tasks are as follows:

`--grid input.h5` Read pressure and wavenumber from this file: output files will use this resolution grid.

`-o|--output output.h5` Write the output to the specified file.

`--scenario str` Add a `scenario` global attribute to the output file.

`--column-range M N` Only process columns `M` to `N`.

To generate the solar spectral irradiance at the spectral resolution of the gases, provide one of the gas spectra files to the `--grid` argument above, and additionally the following arguments:

`--ssi TSI ssi.h5` Create a solar spectral irradiance file with a total solar irradiance of `TSI` (in W m^{-2}).

A suitable file to use here is `data/mean-ssi_nrl2.nc` in the CKDMIP software package, which contains the solar spectral irradiance from Coddington et al. (2016) in units of W m^{-1} cm (i.e. irradiance per unit wavenumber) at up to 1-nm resolution. It is linearly interpolated to the spectral resolution of the gases. The output spectral irradiance is in units of `W m-2`, and sums to the requested `TSI` value. The `work_sw/make_ssi.sh` in the CKDMIP software package demonstrates the use of this feature.

To generate a file of Rayleigh-scattering optical depth at the same spectral resolution as gas absorption, use the following argument:

`--rayleigh` Create a Rayleigh scattering spectral file.

The `work_sw/make_rayleigh.sh` demonstrates the use of this feature.

3.2 Producing reference flux calculations

For this application, `ckdmip_lw` takes as input one or more absorption-spectra files (see section 2.2), optionally scales them to represent different gas concentrations, and sums the layerwise optical depths internally to produce the combined optical depth of all relevant gases. It then performs longwave radiative transfer calculations and outputs profiles of broadband and, optionally, spectral fluxes and fluxes in bands. Note that despite the absorption-spectra files containing temerature, pressure and mole fraction on both full and half levels, `ckdmip_lw` only reads pressure and temperature on half levels and mole fraction on full levels.

The relevant command-line arguments are as follows:

`input.h5` Read in spectral optical depths from the specified file and add them to the optical depth of the gas mixture. Note that this can either be a file for a single gas, or the file containing a merge of several gases (see section 3.3).

`--scale X input.h5` Read the spectral optical depths from the specified file, scale them by a factor `X`, and add the to the optical depth of the gas mixture.

`--conc X input.h5` Read the spectral optical depths and scale them by a constant such that the nominal surface mole fraction (in mol mol^{-1}) is `X`. This is only possible if the file contains a `reference_surface_mole_fraction` variable (i.e. all gases except H_2O and O_3).

`--const X input.h5` Read the spectral optical depths and scale them in each layer so that the mole fraction is `X` in all layers. This is only possible if the file contains a `reference_surface_mole_fraction` variable.

`-c|--config config.nam` Read extra configuration information from the specified namelist file (see below).

`-o|--output output.h5` Write the output to the specified file.

`--scenario str` Add a `scenario` global attribute to the output file.

`--column-range M N` Only process columns `M` to `N`.

Note that multiple input files can be specified. An example of how `ckdmip_lw` was used in this way to generate the files described in section 2.3 may be found in the `work_lw/run_lw_lbl_evaluation.sh` script of the software package. Since the 50 profiles of the Evaluation-1 dataset are stored in five files for each gas, each containing 10 profiles, this script works on 10 profiles at a time, and then uses the NCO utility `ncrcat` to concatenate the profiles into a single file containing the irradiance profiles for all 50 profiles.

If a namelist file is provided then it can be used to set the variable names that are read or written, and other aspects of the processing and output:

```
&longwave_config
pressure_name = "p_hl",           ! Override default, which is "pressure_hl"
pressure_scaling = 100.0,         ! Scaling if provided in hPa (default 1)
temperature_name = "t_hl",        ! Override default, which is "temperature_hl"
optical_depth_name = "od",        ! Override default, which is "optical_depth"
planck_name = "planck",           ! Override default, which is "planck_hl"
wavenumber_name = "wn",           ! Override default, which is "wavenumber"
nspectralstride = 2,              ! Skip spectral intervals (default 1)
do_write_spectral_fluxes = false, ! Write spectral fluxes to output (default no)
do_write_planck = false,          ! Write spectral Planck function (default no)
do_write_optical_depth = false,   ! Write merged optical depth (default no)
band_wavenumber1(1:13) = 0, 350, 500, 630, 700, 820, 980, 1080, 1180, &
  1390, 1480, 1800, 2080,         ! Specify lower wavenumbers of output bands
band_wavenumber2(1:13) = 350, 500, 630, 700, 820, 980, 1080, 1180, 1390, &
  1480, 1800, 2080, 3260,         ! Specify upper wavenumbers of output bands
iverbose = 3,                     ! Set verbosity level from 1 to 5
/
```

Here, the `band_wavenumber1` and `band_wavenumber2` entries specify that in addition to broadband and (optionally) full spectral output, the irradiance profiles should also be provided in bands bounded by the wavenumbers in these two vectors (in cm^{-2}). This is useful to evaluate CKD schemes in individual bands.

3.3 Merging gases

This usage of `ckdmip_lw` allows the optical depth of 'hybrid' or 'composite' gases to be created, by combining the following option with those in section 3.2.

`-m|--merge-only` Store the merged optical depth of the gas mixture in the output file, but do not perform any radiative transfer calculations.

See the scripts `work_lw/test_merge_only.sh` and `test_lw_lbl_merged.sh` in the CKDMIP software package.

3.4 Producing fluxes for CKD models

The `ckdmip_lw` program can also work on files produced by CKD models, by replacing all the options in section 3.2 related to reading in absorption spectra with a single option:

`--ckd input.nc` Read the optical depth and Planck function in each g-point of a CKD model from the specified file.

In addition to broadband fluxes, the output file will include the flux profiles for each of the g-points used by the CKD model.

The idea is that the CKD model produces a description of its approximation of the optical properties of the atmosphere, but does not perform any radiative transfer. By using `ckdmip_lw` to do the radiative transfer, we ensure that the differences with line-by-line fluxes and heating rates (produced by `ckdmip_lw` in section 3.2) are due only to the treatment of gas absorption. The format of the input file is described in section 4. The `ckdmip_lw` program then essentially applies Eqs. 7–12 of Hogan and Bozzo (2018) for the radiative transfer. Note that the surface is assumed to have an emissivity of 1 and the same temperature as the lowest half level of the atmosphere.

Once radiative transfer has been run on the output from a CKD model, the fluxes and heating rates may be compared to those of the line-by-line reference calculations (section 2.3). The Matlab code in the `matlab` directory of the package illustrates how this may be done.

4 Requirements of participants

The purpose of CKDMIP is primarily to compare *CKD tools,* which we define as a method (which may be fully automated or involve some hand-tuning) for generating individual CKD models, with some means to control the trade-off between accuracy and the number of g-points. An individual *CKD model* is a software component for converting profiles of temperature, pressure and gas concentrations into optical depth profiles in N g-points (i.e. to produce the input illustrated in section 3.4), where N is specific to a particular CKD model.

Some CKDMIP participants will have a CKD tool that, with relative ease, can generate CKD models targeted at a particular application and adjust the trade-off between accuracy and efficiency. Section 4.2 of Hogan and Matricardi (2020) describes the CKD models that would ideally be generated by such participants. Other participants will have tools that are less flexible, perhaps involving hand-tuning, so would only be able to generate a small number of models.

Either way, for each generated CKD model, the scenarios that it should be run on are given in section 4.2 of Hogan and Matricardi (2020). Thus for each model–scenario pair, the participant should submit a file containing the optical depth profiles in each g-point that can be used to compute fluxes as described in section 2.3. These files should have names consisting of seven fields separated by underscores:

`TOOL_DATASET_BAND_APPLICATION_CONFIGURATION_optical-depth_SCENARIO.nc`

where the fields are:

`TOOL` A string identifying the CKD tool and possibly its version number, e.g. `ecrad-rrtmg` for the ecRad implementation of the RRTMG gas optics scheme, or `ecmwf-fsck-v0.1` for version 0.1 of an FSCK scheme developed at ECMWF.

`DATASET` Either `evaluation1` or `evaluation2` depending on which set of 50 profiles are being processed.

`BAND` Either `lw` or `sw`.

`APPLICATION` The 'application' from the three listed in Table 1 of Hogan and Matricardi (2020): `limited-area-nwp`, `global-nwp` or `climate`.

`CONFIGURATION` Participants are requested to generate several CKD models to explore the accuracy–efficiency trade-off: this string is used to distinguish between these different configurations, and would typically include the total number of g-points. For example, RRTMG could use `narrow-n140` since its bands fall into the 'narrow bands' shown in Table 4 of Hogan and Matricardi (2020) and it uses a total of 140 g-points, while another scheme could use `wide-n80`.

`SCENARIO` The gas-concentration scenario to which the CKD model has been applied (see Table 2 and section 4.2 of Hogan and Matricardi, 2020), using strings shown in the last field of the file name shown in section 2.3.

The file must be in NetCDF format (either 'classic' or HDF5). Its format is best illustrated by showing the `ncdump` header output on a conforming file produced by ecRad (Hogan and Bozzo, 2018), which uses the RRTMG CKD model of Mlawer et al. (1997) for its gas optics. In the case of the 'Evaluation-1' dataset in present-day conditions we have:

```
netcdf ecrad-rrtmg_evaluation1_lw_climate_narrow-n140_optical-depth_present {
dimensions:
    column = UNLIMITED ; // (50 currently)
    level = 54 ;
    gpoint_lw = 140 ;
    half_level = 55 ;
variables:
    float pressure_hl(column, half_level) ;
        pressure_hl:long_name = "Pressure on half-levels" ;
```

```
        pressure_hl:units = "Pa" ;
    float od_lw(column, level, gpoint_lw) ;
        od_lw:long_name = "Clear-sky longwave optical depth" ;
        od_lw:units = "1" ;
    float planck_hl(column, half_level, gpoint_lw) ;
        planck_hl:long_name = "Planck function on half-levels" ;
        planck_hl:units = "W m-2" ;

// global attributes:
        :title = "Spectral radiative properties from the ecRad radiation model" ;
        :source = "ecRad offline radiation model" ;
...
```

The layers must start at the top-of-atmosphere, but otherwise the format is quite flexible. The dimension names are not important; `ckdmip_lw` only cares about the ordering of dimensions of the three variables. The column-like dimension is shown here as unlimited but this is not a requirement. In this example, the number of longwave g-points is 140, but this number will of course be different for other CKD models. The global attributes are not read.

The three variables shown here provide the information needed to perform radiative transfer by `ckdmip_lw` in section 2.3. They are the pressure and the Planck function at half levels (layer interfaces) and the optical depth at full levels (layers). The Planck function for each g-point is the integral across all parts of the spectrum that contribute to that g-point. The actual names of these three variables can be anything, since the namelist file (see section 3.2) specifies the names of the variables to look for.

One additional file is required for each CKD model, describing what fraction of different parts of the spectrum contribute to each g-point. This will be used as part of the investigation of the effect of clouds described in section 4.4 of Hogan and Matricardi (2020). In the longwave this should be expressed in 326 spectral intervals at a resolution of 10 cm^{-1} between 0 and 3260 cm^{-1}, and in the shortwave it should be provided in 995 spectral intervals at a resolution of 50 cm^{-1} between 250 and 50,000 cm^{-1}. This is commensurate with the spectral scale at which the optical properties of clouds varies, and also ensures that each spectral interval lies entirely within one of the 'narrow' spectral bands defined in Tables 4 and 5 of Hogan and Matricardi (2020). The file name should be of the form

`TOOL_BAND_APPLICATION_CONFIGURATION_spectral-definition.nc`

and an example file for the ecRad implementation of RRTMG is available from the CKDMIP web site, with a `ncdump` header output of:

```
netcdf ecrad-rrtmg_lw_climate_narrow-n140_spectral-definition {
dimensions:
    gpoint_lw = 140 ;
    wavenumber = 326 ;
variables:
    float wavenumber1(wavenumber) ;
        wavenumber1:long_name = "Lower wavenumber bound of spectral interval" ;
        wavenumber1:units = "cm-1" ;
    float wavenumber2(wavenumber) ;
        wavenumber2:long_name = "Upper wavenumber bound of spectral interval" ;
        wavenumber2:units = "cm-1" ;
    float gpoint_fraction(gpoint_lw, wavenumber) ;
        gpoint_fraction:long_name = "Fraction of spectrum contributing to each g-point" ;
```

The variables `wavenumber1` and `wavenumber2` define the bounds of the high-resolution spectral intervals. The variable `gpoint_fraction` the defines what fraction of the spectrum at this high spectral resolution contributes

to each g-point. This variable should sum to unity along the `wavenumber` dimension. Note that in the longwave, the fractions should *not* be weighted by the Planck function at a particular temperature; since the spectral intervals are at high resolution, they can be weighted by the Planck function afterwards, using any temperature desired. Likewise, in the shortwave they should not be weighted by the solar spectrum.

References

Clough, S. A., M. W. Shephard, E. J. Mlawer, J. S. Delamere, M. J. Iacono, K. Cady-Pereira, S. Boukabara and P. D. Brown, 2005: Atmospheric radiative transfer modeling: a summary of the AER codes. *J. Quant. Spectrosc. Radiat. Transfer,* **91,** 233–244.

Coddington, O., J. L. Lean, P. Pilewskie, M. Snow and D. Lindholm, 2016: A solar irradiance climate data record. *Bull. Am. Meteorol. Soc., 97,* 1265–1282.

Hogan, R. J., and Bozzo, A., 2018: A flexible and efficient radiation scheme for the ECMWF model. *J. Adv. Modeling Earth Sys.,* **10,** 1990–2008.

Hogan, R. J., and M. Matricardi, 2020: Evaluating and improving the treatment of gases in radiation schemes: the Correlated K-Distribution Model Intercomparison Project (CKDMIP). To be submitted to *Geosci. Model Dev.,* draft available from https://confluence.ecmwf.int/display/CKDMIP.

Mlawer, E. J., S. J. Taubman, P. D. Brown, M. J. Iacono, and S. A. Clough, 1997: Radiative transfer for inhomogeneous atmospheres: RRTM, a validated correlated-k model for the longwave. *J. Geophys. Res. Atmos.,* **102,** 16 663–16 682.