Impact of different spectral averaging methods in ecCKD

The new ECMWF gas optics tool "ecCKD" first creates a look-up table by averaging the spectral absorption coefficients represented by a single "k term" to a single value of mass absorption coefficient for each gas as a function of temperature, pressure and (in the case of water vapour) concentration. It then performs an optimization step in which the look-up table coefficients are optimized to minimize flux and heating-rate errors against a set of 50 training profiles (the Evaluation-1 CKDMIP dataset). It turns out that this optimization step is crucial for accuracy, but it is nonetheless interesting to see how different averaging methods perform before the optimization step. In the following plots we show the error in terms of present-day fluxes and heating rates evaluated using the "Evaluation-2" CKDMIP dataset. The "narrow" band structure is used, consisting of 13 bands (mapping closely the ones used by RRTMG), and a total of 64 k terms are used.

Linear and logarithmic averaging

The left set of plots below shows the result of linear averaging, the right logarithmic averaging. We see that linear averaging leads to optical depth being overestimated: the atmosphere is too opaque resulting in an underestimate of outgoing longwave radiation (OLR) and an overestimate in surface downwelling. Logarithmic averaging results in a weak overestimate in OLR suggesting the atmosphere is now too optically thin, although the surface downwelling is quite good.



Transmission averaging

By default, ecCKD converts layer optical depths OD to transmittances T for a diffusivity of 1.66, i.e. T=exp(-1.66*OD), averages the transmittances, then converts back to effective layer optical depth. The layering is the same as that of the CKDMIP "Idealized" dataset, which is logarithmic in pressure with ten layers per decade. This approach is dependent on the layer thickness, and it is far from clear what the appropriate effective layer thickness should be. In the limit of very thin layers, the averaging of transmittances would be equivalent to linear averaging of optical depths. In the plots below we multiply the layer optical depths by 1, 2, 3 and 10 before computing the transmittances, which are then averaged spectrally and converted back to optical depths. Thus, the multiplication by ten would be equivalent to computing the transmittance over a layer covering a factor of 10 in pressure.

1x	2x



We see from the results above that the 1x transmission averaging tends to lead to the atmosphere being too opaque (OLR is too low) while 10x transmission averaging leads to the atmosphere being too transparent. 3x transmission averaging gives the least biased OLR while 2x transmission averaging gives the least biased out of the transmission averaging approaches leads to larger errors in heating rates in the very lowest layer of the atmosphere, compared to either the logarithmic or linear averaging.

Optimized solution

For comparison, the plot below shows the results **after** the optimization step (optimizing an original look-up table using 1x transmission-averaging). Clearly the accuracy is very much greater, but there are some outliers in the downwelling fluxes; these are likely profiles from the Evaluation-2 dataset that have conditions (temperature, pressure and/or ozone) that are too far from any of the 50 training profiles for the optimization step to have properly trained the relevant coefficients. Therefore there is still value in ensuring the initial averaging is as accurate as possible.



Results for instantaneous radiative forcing

It is also interesting to look at the performance in terms of radiative forcing, perturbing each of the five major greenhouse gases in turn, shown below. In this case the logarithmic averaging performs quite poorly, especially for N2O. This could be because for the more minor gases, the k terms "specializing" in them are fewer, so span a larger range of absorption; therefore the averaging method plays a more important role.







