**Geoscientific
Model Development**

The Lagrangian chemistry and transport model ATLAS: validation of advective transport and mixing

I. Wohltmann and M. Rex

Alfred Wegener Institute for Polar and Marine Research, Potsdam, Germany

Abstract. We present a new global Chemical Transport Model (CTM) with full stratospheric chemistry and Lagrangian transport and mixing called ATLAS (Alfred Wegener InsTitute LAgrangian Chemistry/Transport System). Lagrangian (trajectory-based) models have several important advantages over conventional Eulerian (grid-based) models, including the absence of spurious numerical diffusion, efficient code parallelization and no limitation of the largest time step by the Courant-Friedrichs-Lewy criterion. The basic concept of transport and mixing is similar to the approach in the commonly used CLaMS model. Several aspects of the model are different from CLaMS and are introduced and validated here, including a different mixing algorithm for lower resolutions which is less diffusive and agrees better with observations with the same mixing parameters. In addition, values for the vertical and horizontal stratospheric bulk diffusion coefficients are inferred and compared to other studies. This work focusses on the description of the dynamical part of the model and the validation of the mixing algorithm. The chemistry module, which contains 49 species, 170 reactions and a detailed treatment of heterogeneous chemistry, will be presented in a separate paper.

1 Introduction

Chemical Transport Models are commonly based on a fixed spatial grid and include an Eulerian advection scheme for the transport of chemical species (e.g. SLIMCAT, Chipperfield, 2006). Although much effort has been put into the development of advection schemes that minimize numerical diffusion (e.g. Prather, 1986), numerical diffusion is in principle unavoidable in Eulerian models.

In the context of this paper, numerical diffusion is defined as any process that is triggered by the model calculations and behaves like a diffusive process acting on the chemical species in the model. Typically, numerical diffusion will be spurious and unrealistic, e.g. present in every grid cell with the same magnitude. We do not consider processes as numerical diffusion here which are deliberately introduced into the model, e.g. to mimic the observed diffusion. Eulerian approaches suffer from numerical diffusion because they contain an advection step that transfers information between different grid cells. Usually, this step is connected to averaging over several grid points, which introduces diffusion on the spatial scale of the grid resolution at the temporal scale of the model time step. An example is the interpolation of species to trajectory end points in Semi-Lagrangian approaches.

Since diffusion is very small in the stratosphere, numerical diffusion in Eulerian models is always much larger than the real atmospheric diffusion (e.g. Konopka et al., 2005), leading to problems like excessive mixing across atmospheric mixing barriers. Particularly at the low resolutions typically required for long runs (e.g. where the CTM module is coupled to a General Circulation Model, GCM, as part of a Chemistry Climate Model, CCM), numerical diffusion can lead to problems in the representation of the polar vortices and hence the Antarctic ozone hole. With an Eulerian model it is difficult to maintain the steep gradients of species at the boundary of the ozone hole, which are not only relevant for a proper representation of ozone depletion, but also for the radiative and dynamical feedback from the ozone hole on climate processes. We have therefore developed a Lagrangian CTM that can be used as a stand alone model or as module in a CCM system.

Lagrangian models have several advantages over Eulerian models:

- There is no numerical diffusion. Actually, some explicit form of diffusion has to be included in Lagrangian models to obtain realistic results, which gives

the opportunity to model the real atmospheric diffusion on a physical basis.

– The model architecture allows for an efficient parallelization of code and runs fast on low-cost computers.

– The time step of the integration of the transport equations can be chosen freely. There is no limitation of the largest time step by the Courant-Friedrichs-Lewy criterion.

– Mixing ratios of inert tracers are conserved by design (no negative mixing ratios by transport).

– Transport of additional species does not add to computing time.

In addition, our specific model is free of any grid in the vertical or horizontal. The absence of numerical diffusion, the conservation of mixing ratios and the fact that transport of additional species does not add to computing time are basically due to the fact that the Lagrangian time derivative of mixing ratios is zero in the absence of local sources and sinks, while the Eulerian derivative contains an advection term. Despite these obvious advantages, only very few global Lagrangian chemical transport models with explicit mixing exist so far (e.g. Collins et al., 1997; Fairlie et al., 1999). The most prominent examples are the CLaMS model (McKenna et al., 2002; Konopka et al., 2004, 2007) and the ATTILA model (Reithmeier and Sausen, 2002), which has recently been coupled to the General Circulation Model ECHAM (Stenke et al., 2009).

The main focus of this model version is on a proper representation of the chemistry and transport of the stratosphere, although a basic troposphere is included. Hence, the validation focusses on the stratosphere. It is planned to extend the model to the troposphere in the future.

Both resolution and trajectory time step can be chosen freely, so the model can be used both for highly resolved process studies and long-term runs over several decades. One of the specific advantages of Lagrangian models is that they allow to study the fine filamentary structures that are usually below the resolution of conventional Eulerian models and that disappear rapidly due to numerical diffusion. This is possible since the information over the small scale behaviour of tracers is contained in the large-scale wind fields (Batchelor regime in fluid dynamics, see Haynes and Vanneste, 2004).

The main concept of the model is similar to the CLaMS 3-D model (Konopka et al., 2004), which has a more sophisticated mixing parameterization compared to models like e.g. ATILLA or Collins et al. (1997). However, ATLAS is a completely new model developed from scratch, which has no code in common with CLaMS. Some changes and improvements are introduced to the transport and mixing approach of CLaMS here, and their impact on the results is elucidated. The method of validation presented in Konopka et al. (2004)

is also analyzed, and some changes are proposed for a more robust calculation of the validation parameters. Finally, the model is used to infer values for the effective horizontal and vertical diffusion coefficients for the stratosphere, and results are compared with other studies.

Section 2 contains the basic model description. Section 2.1 gives an overview of the model concept. Sections 2.2 and 2.3 describe the model implementation. While Sect. 2.3 gives a detailed treatment of the mixing algorithm, Sect. 2.2 describes the other relevant features of the model. The validation of the mixing algorithm is presented in Sect. 3. Results for diffusion coefficients are discussed in Sect. 4. Conclusions are given in Sect. 5.

2 Model description

2.1 Model concept

The model design is based upon a modular approach. This allows a flexible configuration and easy exchange of components. The basic layout of the model is shown in Fig. 1.

Transport and chemistry in the model are driven by external meteorological fields. These can be provided either by a General Circulation Model (GCM module in the sketch, red box) or by existing analysis data like ECMWF ERA Interim (Simmons et al., 2006).

With this input, a large number of trajectories, each symbolizing an air parcel, is initialized and advected for some hours (filling the domain of the complete atmosphere) (trajectory module, green box). Chemistry is calculated on every trajectory like in a box model (chemistry module, yellow box).

Since there is no numerical diffusion in a Lagrangian model, some form of mixing has to be introduced to obtain realistic trace gas distributions. This has the advantage that mixing can be modeled from physical considerations and that the strength of mixing can be adjusted to observed values, while numerical diffusion cannot be controlled in strength and location. Every few hours, a mixing step is introduced (mixing module, orange box). In this step, new air parcels are added and existing air parcels are merged in regions of large flow deformation (shear and strain) to keep the density of air parcels constant. The actual mixing is accomplished by averaging the chemical species on the new or merged air parcels. As a side effect, adding and merging parcels also ensures that voids and crowded regions are avoided. After the mixing step, the model continues with the advection step again. A detailed justification of the validity of the mixing approach can be found in Sect. 2 of McKenna et al. (2002) and, shortly recapitulated, in Sect. 2.3.2.

2.2 Basic model features

The initial air parcels are distributed randomly over the model domain by drawing uniformly distributed random
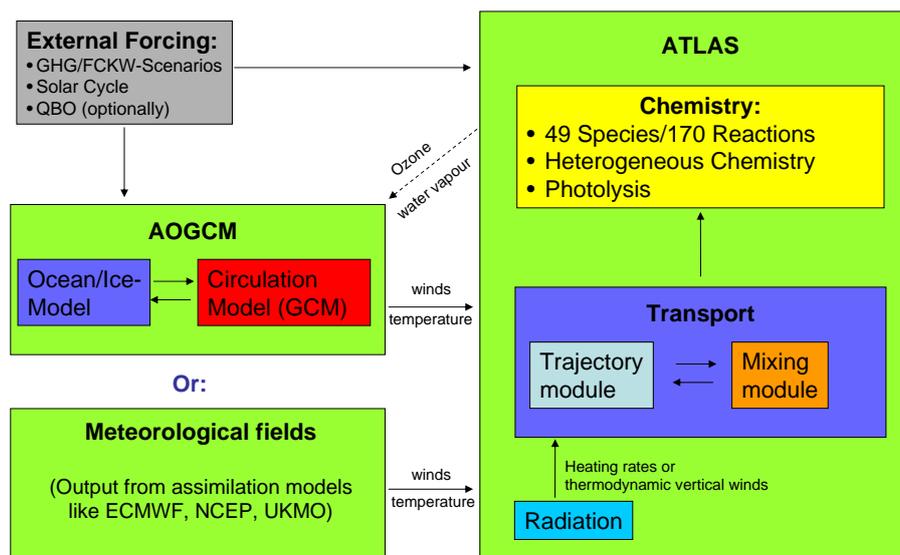
Fig. 1. Model concept.

numbers for longitude, the cosine of latitude and the vertical coordinate. The initial value of the random generator can be set to a fixed value to obtain reproducible results.

Every air parcel contains variables for a configurable number of chemically active species and additionally, for a configurable number of chemically inert tracers.

The vertical coordinate is limited by an upper and lower boundary, which can be chosen freely. The number of air parcels is then chosen such that a predefined initial resolution r_0 is obtained (defined by the mean horizontal parcel distance in layers of the vertical depth Δz introduced in Sect. 2.3.3). While it is possible to set the lower boundary to the surface, the model does only contain a basic representation of the troposphere, mainly used to provide the stratosphere with tropospheric source gases. Several components that would be necessary for a realistic representation of the troposphere are currently not implemented, like tropospheric chemistry, clouds or a treatment of sub-scale convection.

The trajectory module is based on the parallel trajectory code used in Wohltmann and Rex (2008). Time step and integration method (e.g. 4th order Runge-Kutta) are freely configurable. The vertical coordinate system and the associated vertical winds can be set to four different options:

- Pressure as vertical coordinate and vertical winds taken from the input files of the analysis data or GCM, which usually means winds that satisfy mass conservation by the continuity equation.

- Pressure and vertical winds from the thermodynamic equation (satisfying conservation of thermal energy) as in Wohltmann and Rex (2008).

- Potential temperature and heating rates (satisfying conservation of thermal energy).

- A hybrid coordinate that transforms from a pressure coordinate at the surface to a potential temperature coordinate in the stratosphere, as defined in Konopka et al. (2007). The hybrid winds transform from analysis winds in pressure coordinates to heating rates.

The model contains a lower and an upper boundary layer. The depth of the boundary layers is again based on Δz (Sect. 2.3.3). All parcels in these boundary layers are deleted in every model step and replaced by newly initialized parcels. The reinitialization of the parcels is necessary to avoid that holes or volumes with parcels outside the model boundary appear due to the large-scale motion. Parcels outside the model domain are deleted. Tracers and chemical species are interpolated from the old parcel positions to the new parcels where possible. Where not possible, tracers and chemical species are set to climatological values. This may e.g. be the case where the downward branch of the Brewer-Dobson circulation produces volumes void of any parcels directly below the upper boundary.

A tropopause is needed for several purposes in the model. It is set to the 2 PVU potential vorticity surface in the extratropics (dynamical tropopause) and to the 375 K potential temperature surface in the tropics.

The model code is platform independent and can run in parallel on a simple personal computer network connected via the internet. Main modules are implemented in Matlab and time critical code in C. Extensive use of free software libraries is made, the most important ones are Qhull for the triangulation, KPP (Kinetic Preprocessor) for the chemistry module and a Matlab implementation of a parallel MPI like interface based on exchange via ssh.

2.3 Mixing algorithm

The basic idea of the mixing algorithm is identical to the approach used in the CLaMS 3-D model (Konopka et al., 2004). The algorithm is as follows:

- For every air parcel, determine all air parcels that are, in a certain sense (see Sect. 2.3.1 below), next neighbors of that air parcel.

- Advect air parcels with trajectory model and optionally calculate chemistry.

- If the distance between a particular air parcel and one of its former next neighbors (prior to advection) exceeds a critical distance r_+ after advection, insert a new air parcel in the middle between the original air parcel and the neighbor, the new mixing ratio is the mean of the mixing ratios of the air parcel and its neighbor. Do that for every neighbor.

- Determine next neighbors again.

- If the distance between a particular air parcel and several of its neighbors is lower than a critical distance r_-, merge all of these neighbors and the air parcel itself to a new air parcel, the new mixing ratio is the mean of the old mixing ratios. All merged air parcels are removed immediately and can only contribute to one new air parcel.

- Advance to next mixing step.

2.3.1 Next neighbors

Next neighbors are determined by Delaunay triangulation: the air parcels are connected by a mesh of triangles (2-D case) or tetrahedrons (3-D case), such that the circumcircles of any triangle (or circumspheres of any tetrahedron) are empty (e.g. O'Rourke, 1998). Air parcels that are directly connected in the mesh are considered next neighbors. Depending on the model resolution, we use two different methods, which are outlined in Fig. 2:

- Global layer approach (CLaMS): for high model resolutions ($r_0 < 150$ km), the algorithm also applied in CLaMS 3-D (Konopka et al., 2004) is used. The atmosphere is divided into a set of layers. In each layer, the vertical coordinate of the air parcels is ignored and a global 2-D triangulation on the surface of a sphere is performed to obtain the next neighbors. The vertical depth Δz of the layers is determined with one of the methods described in Sect. 2.3.3 to obtain realistic vertical diffusion coefficients and mixing (with z as the log-pressure altitude for the pressure coordinate and appropriate definitions for the potential temperature and hybrid coordinates). A staggered set of layers is used in every second mixing step to avoid a clustering of the air

parcels in the vertical direction by the averaging of the vertical coordinate in new and merged air parcels (see Sect. 2.3.4).

- Local layer approach (ATLAS): for low model resolutions ($r_0 \geq 150$ km), a new method is introduced. The main difference to the global layer approach is that we use a local layer with a limited horizontal extent for every air parcel, which is vertically centered on every single air parcel. The advantage of this approach is that it is less diffusive and gives superior results with the same mixing parameters (see Sect. 3.5). Conceptually, it is also a more realistic approximation to atmospheric mixing, since air parcels that are vertically close, but in adjacent layers do not mix in the CLaMS approach in contrast to the ATLAS approach. Unfortunately, the approach is computationally too expensive for very high resolutions.

As in the global layer approach, the air parcels tend to cluster vertically in the local layer approach. We add a random component to the vertical coordinate of newly inserted or merged air parcels to avoid the generation of artificial layers in the local layer approach. The addition of the random component is explained in more detail in Sect. 2.3.4.

In the first step of the method, a global 3-D triangulation is carried out to get a preliminary set of next neighbors for every air parcel (for some details, see Appendix A). The list of next neighbors is then extended by the second neighbors (next neighbors of next neighbors) to get a cloud of air parcels around every original air parcel (the obvious approach to include all air parcels below a given distance to the original parcel in the cloud is computationally to expensive).

Now, all air parcels outside an interval $\left[z - \Delta z_{\mathrm{low}}, z + \Delta z_{\mathrm{upp}}\right]$ are deleted from the list of neighbors to center the air parcels in the local layer. The vertical mixing depth $\Delta z = \Delta z_{\mathrm{low}} + \Delta z_{\mathrm{upp}}$ is determined as described in Sect. 2.3.3. The remaining parcels and the original air parcel are now treated with a local 2-D triangulation ignoring the vertical coordinate to get a final list of next neighbors. In contrast to the global layer approach, a new 2-D triangulation is necessary for every single air parcel, but the number of points in each triangulation is much lower.

In principle, the final next neighbors could also be obtained directly from the 3-D triangulation. However, there is no constraint on the minimum vertical distance in this approach: newly inserted air parcels will tend to introduce new layers of air parcels with a lower mean vertical distance between parcels than in the last mixing step. These parcels would shield the old parcels that are a little bit further away in the 3-D triangulation and cause a finer mesh in the vertical. This
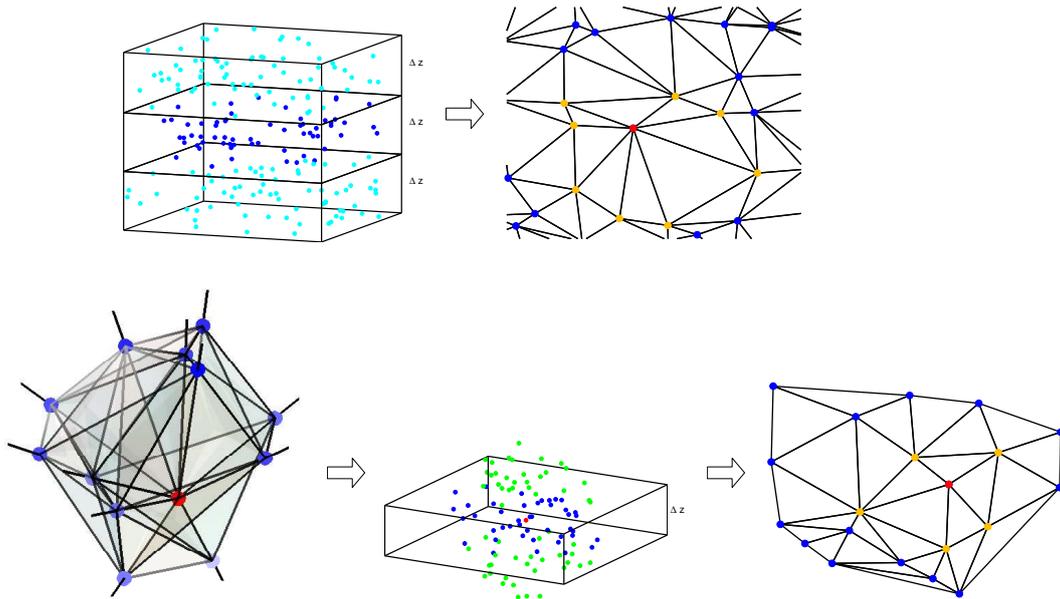
Fig. 2. Sketch of the two methods for finding next neighbors. Top: in the CLaMS approach the atmosphere is divided into layers (left, blue and cyan dots denote the air parcels). In every layer, the z coordinate is ignored and a global 2-D triangulation is performed to find the next neighbors (right, view from above, in the example the orange parcels are the next neighbors of the red parcel, clipped lines symbolize the global triangulation without boundaries). Bottom: in the ATLAS approach, a global 3-D triangulation is performed (left, clipped lines again symbolize the global triangulation) and in a first step, next and second neighbors are determined for each air parcel (middle, in the example all parcels are next or second neighbors of the red parcel). All parcels outside a vertical depth Δz are sorted out (green parcels). The remaining blue parcels are triangulated locally in 2-D ignoring the z coordinate (right).

problem will get worse in every mixing step. In the 2-D triangulation, all these air parcels will tend to end up in the same layer and increase the density of parcels and hence, the number of parcels below the critical distance r_- for merging.

Note that there is no fixed model grid at all in the ATLAS triangulation approach. There are neither distinct layers nor a horizontal grid in the model. The (virtual) model layers in the following discussion just mean all points within a certain interval $[z-\Delta z_{low}, z+\Delta z_{upp}]$ with z chosen freely and the vertical extent chosen to get an appropriate density of points in the horizontal. This is done in the post-processing of the model results and completely independent from the model parameterization.

2.3.2 Horizontal parameterization of mixing

Since the stratosphere is stably stratified, vertical turbulence is small and it is mainly the two-dimensional horizontal chaotic advection that will produce fine horizontal filaments and fractal structures in tracers on different scales. In conjunction with the vertical wind shear, sloping structures with large horizontal-to-vertical aspect ratios will be created (Haynes and Anglade, 1997). These structures will get smaller by the rapid scale collapse induced by the large-scale flow (but maintain their aspect ratio) and finally be dissipated by molecular diffusion and with the help of small-scale turbulence (e.g. by breaking gravity waves), which mixes the

structures down to the molecular scale (e.g. McKenna et al., 2002; Legras et al., 2005; Balluch and Haynes, 1997; Wilson, 2004). This turbulence is intermittent in space and time due to the stable nature of the stratosphere.

The mixing concept is based on the assumption that mixing by small-scale turbulence predominantly occurs in regions of large flow deformation, e.g. in regions of large vertical shear or horizontal strain. The basic idea is to use the Lyapunov exponent to determine the critical distances r_+ and r_-. It can be shown that large values of the Lyapunov exponent are directly related to large shear or strain rates (McKenna et al., 2002).

The Lyapunov exponent λ is a measure for the rate of separation of trajectories that are initially in close neighborhood. In a simplified sense, it can be defined for trajectories by

$$\mathrm{d}r = \lambda r \, \mathrm{d}t \tag{1}$$

where r is the initial separation of the trajectories at $t=0$ and $r+\mathrm{d}r$ is the separation of the trajectories at $t=\mathrm{d}t$. If the flow exceeds a critical value λ_c of the Lyapunov exponent, we will insert a point. If integrating over a time Δt with constant λ_c and initial separation r_0, this corresponds to using a critical distance $r_+=r_0 \exp(\lambda_c \Delta t)$. λ_c is a free parameter of the model and has to be optimized by comparison of model results with observations. The same is true for the length of the mixing time step Δt. r_0 determines the model resolution

and can be chosen freely (it is identical to the r_0 used in the initialization). To relate r_- to r_+, we note that for incompressible flow, a small circle of radius r_0 will be deformed into an ellipse with the same area after a sufficiently small time Δt. If the minor and major axes are r_- and r_+, that means $r_0^2 = r_+ r_-$ and we get

$$r_\pm = r_0 \exp(\pm \lambda_c \Delta t) \tag{2}$$

avoiding another free parameter in the model. In a typical shear flow, r_+ will apply to diffusion along the trajectory and r_- to diffusion orthogonal to it.

The mixing concept ensures that mixing in the model predominantly takes place where the flow deformation is large. This is in broad agreement with observed mixing barriers like the polar vortex edge and observed well-mixed regions like the surf zone, which are indeed situated where deformation is small or large, respectively (e.g. Haynes and Shuckburgh, 2000).

λ_c controls the intensity of mixing (in terms of the percentage of mixed parcels per time interval). The lower λ_c is set, the more air parcels are affected by mixing. In conjunction with r_0, which specifies the horizontal distance affected by one mixing event, λ_c determines the horizontal diffusion coefficient K_h (for an explicit equation, see Sect. 2.3.5). Hence, λ_c is a crucial parameter to adapt the strength of mixing in the model to observed values.

The parameter r_0 does not need to be the actual resolution r_0^{eff} of the model, since there is no constraint on the total number of parcels in the model. While the initialization starts with a mean parcel distance of r_0, the number of parcels will change as long as the number of inserted parcels and the number of deleted parcels do not agree. Eventually, an equilibrium will be reached where this is the case. The equilibrium value for the effective resolution (defined by the mean parcel distance) is close to r_0 both for small values of λ_c (high mixing intensity leads to high grid adaption frequency) and high values of λ_c (no mixing cannot alter the number of parcels), but lower inbetween (see Fig. 3).

Reasonable values of Δt are limited to a range from approximately 6 to 24 h. For too short mixing times, the mixing step gets computationally too expensive. For too long mixing times, neighbor relationships are not conserved and the non-linearity of the flow prevents the application of the mixing algorithm. Δt is no crucial parameter in the parameterization, since validation results only weakly depend on the product $\lambda_c \Delta t$.

2.3.3 Vertical diffusion coefficient

The actual magnitude of vertical diffusion enters into the model through the depth Δz of the layer mixed in one mixing time step, which is directly related to the vertical diffusion coefficient: the greater the diffusion coefficient, the broader the layer mixed in a given time period. Turbulent
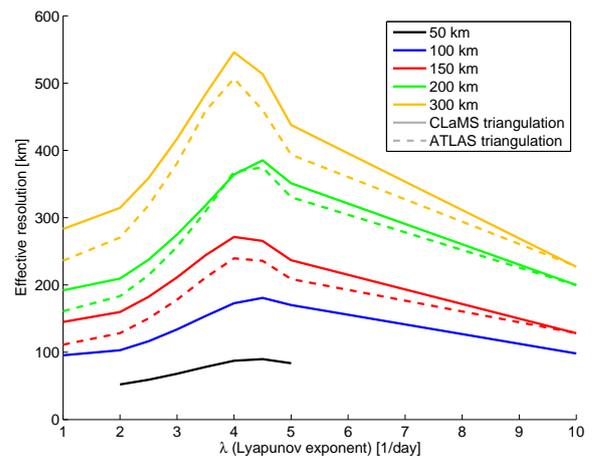

Fig. 3. Effective resolution for different values of the critical Lyapunov exponent λ_c (axis) and resolution r_0 (colors).

diffusion in the stratosphere is dominated by vertical diffusion processes, since the horizontal-to-vertical aspect ratio α of tracer structures in the stratosphere is large (about 250) (Haynes and Anglade, 1997). Tracer filaments will be mixed to background values in vertical direction over the whole horizontal extent of the filament quickly, while in the same time period, horizontal diffusion will only affect the outer edges of the filament. For this reason, the apparent (effective) horizontal diffusion is much larger than the actual horizontal diffusion and directly dependent on the vertical diffusion and aspect ratio by $\alpha^2 = K_h / K_z$.

The standard deviation σ of the vertical positions of air parcels that start at the same vertical coordinate z after a time Δt is related to the vertical diffusion coefficient by

$$K_z = \frac{\sigma^2}{2\Delta t} \tag{3}$$

This and a normal distribution of the air parcels follows from Fick's second law with a delta function as initial condition. In this sense, if $K_z = \Delta z^2 / (8\Delta t)$, a layer of approximately $\Delta z = 2\sigma$ (containing 69% of the air parcels) is mixed after Δt, which gives the possibility to estimate the mixing layer depth from the diffusion coefficient.

Note that the log-pressure altitude $z = H \log(p_0 / p)$, H scale height, is used in all derivations in this section, which do directly apply to the pressure coordinate system. For the potential temperature and hybrid coordinates, Δz is transformed into appropriate quantities in a last step.

A first guess of the vertical diffusion coefficient is obtained from a climatology K_z^{clim} that is a function of log-pressure height (Fig. 4). We denote the diffusion coefficient as K_z^{clim} here to discriminate it from the effective vertical diffusion coefficient K_z^{eff} that is actually applied in the model (see the following paragraphs). The climatology below 50 km is taken from Massie and Hunten (1981), while the part between 50 km and 80 km is following an exponential increase

(e.g. Shimazaki and Wuebbles, 1973). In the troposphere, the coefficient is set to a constant value of $10 \, \mathrm{m^2/s}$. The boundary between the troposphere and the stratosphere is determined by the altitude of the tropopause. In the stratosphere below 15 km, the coefficient is set to a constant value of $0.58 \, \mathrm{m^2/s}$. There is considerable uncertainty in the values of the vertical diffusion coefficient up to the present day (e.g. Wilson, 2004; Legras et al., 2005), and our parameterization, which is based on some rather old, but easily available data, fits well into the published range of values. Note also that the climatology was derived from tracer data by neglecting advective transport, which may introduce uncertainty. Since the climatology is scaled in the following, only the shape and not the absolute values are of importance anyway.

Alternatively, the vertical diffusion coefficient can be set proportional to entropy (the approach taken by CLaMS, Konopka et al., 2007) or to a constant value. Actually, the shape of the diffusion profile based on entropy is quite similar to our approach (Fig. 4, red line), which makes the choice of the method a matter of philosophy rather than being of practical implications.

The vertical mixing depth Δz is obtained individually for each air parcel in the ATLAS triangulation approach. Ideally, Δz would be obtained from

$$\Delta z(z) = \sqrt{8 K_z^{\mathrm{clim}}(z) \Delta t} \tag{4}$$

where z is the log-pressure altitude of the air parcel. However, the equation actually used is

$$\Delta z(z) = \frac{2 r_0}{\alpha(20 \mathrm{km})} \sqrt{\frac{K_z^{\mathrm{clim}}(z)}{K_z^{\mathrm{clim}}(20 \mathrm{km})}} \tag{5}$$

$\alpha(20 \mathrm{km})$ is another free parameter of the model and defines the aspect ratio between the vertical scale Δz and the horizontal scale r_0 of the model in the lower stratosphere. This follows from the above equation with $z{=}20 \, \mathrm{km}$:

$$\alpha(20 \mathrm{km}) = 2 r_0 / \Delta z(20 \mathrm{km}) \tag{6}$$

The factor 2 is due to the fact that in a layer of depth Δz with uniformly distributed points, the mean distance of any two points will be $\Delta z/2$.

The reason for skipping the (physically correct) dependency on Δt and introducing the dependency on r_0 is that it is important to match the horizontal and the vertical scale of the model. The mean horizontal distance between air parcels in a layer of depth Δz must be of the order $r_0(z){=}\alpha(z)\Delta z(z)/2$ to ensure that mixing events affecting a certain distance in the vertical do affect an area of the correct horizontal extent and vice versa, as given by α. A simple example may illustrate that: a typical filament in the stratosphere could have a thickness of 1 km and a horizontal extent of 250 km, corresponding to an aspect ratio of 250. If the vertical diffusion coefficient would mix a layer of 1 km in 12 h to background values, it would also mix the structure of 250 km extent horizontally to background values in that time. A filament of
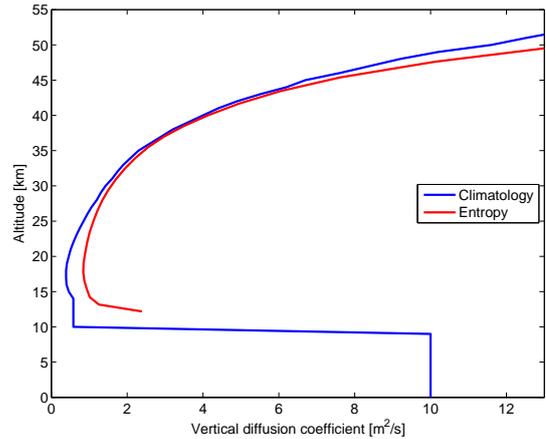

Fig. 4. Vertical diffusion coefficient climatology.

thickness 2 km needs to have a width of 500 km and will need correspondingly longer to be mixed. If the horizontal and vertical resolution of the model would be decoupled, the mixing in the model would not work anymore in a realistic way.

If Eq. (4) would be used for the vertical resolution and the aspect ratio would still be respected, the horizontal resolution would be no free parameter and e.g. $r_0(20 \mathrm{km}){=}\alpha(20 \mathrm{km})\Delta z(20 \mathrm{km})/2$. With Δz as small as inferred from the vertical diffusion coefficient climatology (e.g. $\Delta z {\approx} 450 \, \mathrm{m}$ and $r_0 {\approx} 55 \, \mathrm{km}$ for $K_z {=} 0.58 \, \mathrm{m^2/s}$ and $\Delta t {=} 12 \, \mathrm{h}$), the number of air parcels and model layers would increase to values computationally too expensive. That is, the method used in the model overestimates the effective vertical diffusion coefficient due to computational constraints. Note that CLaMS effectively implements the layer thickness in an equivalent way and suffers from the same problem. This can be partly compensated by setting λ_c to high values, which reduces the number of mixing events. However, while this will lower the effective bulk diffusivity of the model, every single mixing event will still take place over a vertical distance that is too large.

Finally, Δz is split into an upper part Δz_{upp}, which we use to search for all points between z and $z{+}\Delta z_{\mathrm{upp}}$ and a lower part Δz_{low} to search for points between z and $z{-}\Delta z_{\mathrm{low}}$. The simplest approach would be to set both values to $\Delta z/2$. In our approach, Δz_{low} and Δz_{upp} can take different values. First, if the interval $[z{-}\Delta z(z)/2, z]$ crosses the tropopause, Δz_{low} is set such that the interval ends at the tropopause. The same is done with Δz_{upp} and the interval $[z, z{+}\Delta z(z)/2]$. To avoid problems in other regions where K_z changes rapidly with altitude, Δz_{low} is calculated as half of the mean value of $\Delta z(z')$ between z and $z{-}\Delta z(z)/2$,

$$\Delta z_{\mathrm{low}}(z) = \frac{1}{2} \frac{\int_{z-\Delta z(z)/2}^{z} \Delta z(z') \, dz'}{\Delta z(z)/2} \tag{7}$$

An analoguous aproach is taken with Δz_{upp}.

Equation (5) can be interpreted as introducing an effective $\alpha(z)$ by

$$\Delta z(z) = 2r_0/\alpha(z)$$

and

$$\alpha(z) = \alpha(20\,\text{km})\sqrt{K_z^{\text{clim}}(20\,\text{km})/K_z^{\text{clim}}(z)}.$$

This implies a higher effective aspect ratio at altitudes with a lower Δz and vice versa. Very probably, this ad hoc assumption of the variation of $\alpha(z)$ with altitude is not compatible with the real atmospheric variation, but this variation follows from the horizontal resolution of the model held constant with altitude. In future versions, that may be changed to allow a varying horizontal resolution. For the following validation the effect is small, since the model is compared to measurements around 20 km altitude.

2.3.4   Inserting new parcels

If a new air parcel is inserted, its coordinates are calculated as follows: the longitude and latitude of the new air parcel are calculated by transforming all longitudes and latitudes of the old air parcels to cartesian coordinates with a local orthographic projection, and averaging the x and y coordinates. The result is transformed back.

In the CLaMS approach, the vertical coordinate is the average of the old vertical coordinates. This would lead to a clustering of the air parcels in the middle of each layer with fixed layers: the average of some random numbers from an interval will typically be close to the middle of the interval and is bounded by the original numbers. This increases the number of air parcels near the middle of the layer in every mixing step and leads to a positive feedback which results in the generation of layers in the vertical distribution of parcels. For this reason, a staggered set of vertical layers is used in every second mixing step in the CLaMS global layer approach. However, this approach still cannot avoid a certain clustering of the parcels into layers (Fig. 5).

Somewhat surprisingly, there is the same problem in the ATLAS local layer approach of triangulation, if the CLaMS method of averaging is used. Although there are no predefined layers, the same process can be triggered once there is a small random initial deviation from a uniform vertical distribution of the air parcels: the deviation will tend to be amplified in the next mixing steps by the feedback described above. This results in a process of self-organisation of the vertical coordinate values, where neighboring layers of enhanced parcel density are in concurrence for the parcels inbetween.

It is necessary to work against this effect to avoid the formation of pronounced vertical layers and associated problems with vertical mixing. Since there is no possibility of a staggered set of layers as in the CLaMS approach, a random number is added to the average of the vertical coordinates in

every newly inserted or merged air parcel to achieve a more uniform vertical distribution of air parcels (Fig. 5). The random number is chosen from an interval $\left[-\Delta z_{\text{low}}, \Delta z_{\text{upp}}\right]$, with these values again calculated as in Sect. 2.3.3. In most cases, the addition of the random component only slightly increases the diffusivity.

Since there is still a certain vertical clustering of the air parcels in the CLaMS global layer triangulation approach with simple vertical averaging, we also add a random number to the vertical coordinate in the CLaMS triangulation approach by default. This modified CLaMS approach is the only computationally feasible method for higher resolutions, if we want to achieve a distribution as uniform as possible also here. A discussion of the differences of the approaches in the results is given in Sect. 3.5.

2.3.5   The effective large-scale diffusion in the model

It is possible to estimate the diffusion coefficient from the mean distance Δr of any two points in a Gaussian distribution of air parcels as the one used for Eq. (3). It turns out that $\Delta r^2 = 2\sigma^2$ and hence the estimate is

$$K = \frac{\Delta r^2}{4\Delta t} \tag{8}$$

According to this equation, we estimate the diffusion coefficient in the model for every newly inserted or merged parcel. For comparison with other published values of the diffusion coefficient, the diffusion coefficient averaged over a certain geographical area and not for an individual air parcel is needed. Hence, a mean diffusion coefficient is calculated for every mixing step by averaging over the diffusion coefficients of the air parcels in a given area. Parcels not affected by mixing are assigned a coefficient of zero.

It is possible to derive equations that provide some insight on how the diffusion behaves if we change the free model parameters r_0, λ_c, α and Δt, see also Fig. 5 of Konopka et al. (2003). If two air parcels are mixed in the r_+ or r_- step, the horizontal diffusion that the newly inserted air parcel experiences can be roughly estimated to

$$K_{\text{h}\pm}^{\text{eff}} \approx \frac{r_{\pm}^2}{4\Delta t} = \frac{r_0^2}{4\Delta t}\exp(\pm 2\lambda_c\Delta t) \tag{9}$$

and the vertical diffusion to

$$K_z^{\text{eff}} \approx \frac{(\Delta z/2)^2}{4\Delta t} = \frac{r_0^2 K_z^{\text{clim}}(z)}{4\alpha(20\,\text{km})^2 K_z^{\text{clim}}(20\,\text{km})\Delta t} \tag{10}$$

since the typical distances between the air parcels that contribute to the new parcel are $r_{\pm}$ and $\Delta z/2$. It is obvious that the diffusivity of the model increases for decreasing resolution, just because every single mixing event occurs over a larger distance. For decreasing λ_c (more mixing), $K_{\text{h}+}^{\text{eff}}$ (diffusion in flow direction) decreases, while $K_{\text{h}-}^{\text{eff}}$ (diffusion orthogonal to flow direction) increases. For the spatially averaged diffusion and also for the average diffusion that a single
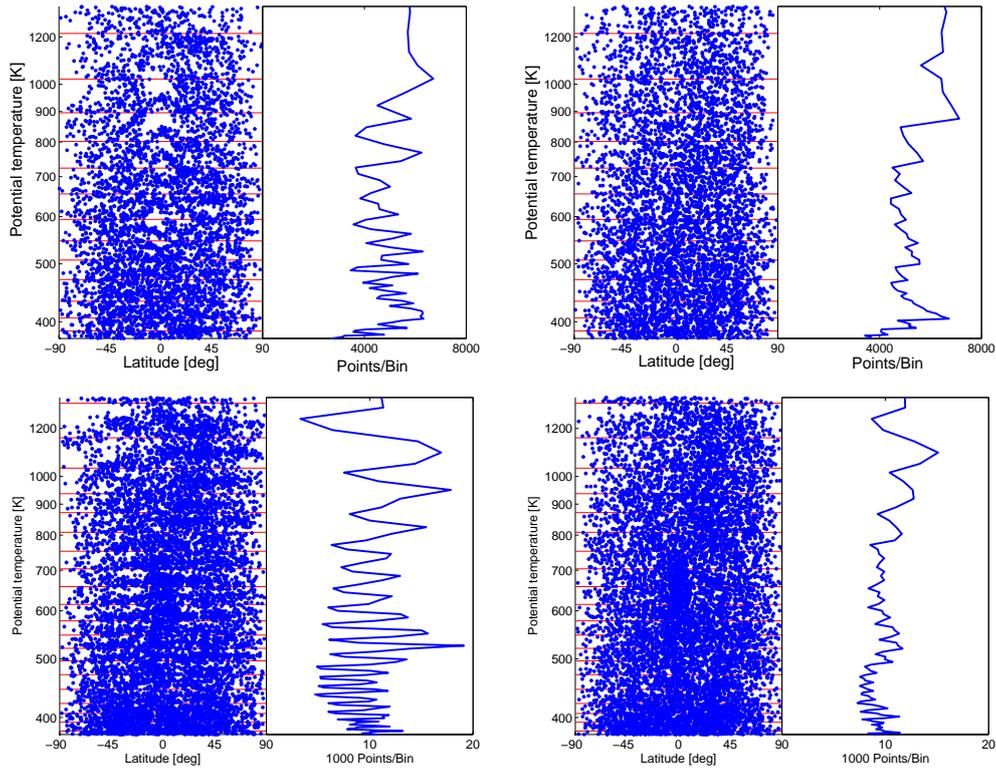
Fig. 5. Vertical distribution of air parcels for the ATLAS triangulation approach (top) or the CLaMS triangulation approach (bottom) and the random vertical coordinate off (left) or on (right). Results of a model run with r_0=150 km (top) or r_0=100 km (bottom), λ_c=2 d^{-1} after 4.5 months. The scatter plot shows the distribution of air parcels in the potential temperature-latitude plane (ignoring longitude). Only a subset of all air parcels is shown to make the clustering visible. The blue line shows the vertical density of the points (as the number of points per the local value of the layer depth Δz). Red lines are the layers of the CLaMS approach.

parcel experiences over a large number of mixing steps, the percentages of added and merged parcels also matters (see Fig. 11). These are also a function of the mixing parameters, particularly of λ_c, where lower values correspond to higher percentages of mixed parcels and a higher average diffusion coefficient (and vice versa). Supplementary to the analysis in Konopka et al. (2003), we find that the increase of diffusivity with lower λ_c in CLaMS and ATLAS is not only due to the increase in K_{h-}^{eff}, but that the higher percentage of mixed parcels contributes significantly. Note that the vertical and horizontal diffusion can be varied independently from each other within certain limits, since λ_c appears only in the horizontal diffusion coefficients and α only in the vertical.

2.3.6 Additional patchyness from noise in the wind fields

In principle the strength of mixing between air parcels can be tuned to the real atmospheric mixing. But in practice random motions from noise in the driving wind fields and their Eulerian discretization result in spurious random transport of air masses. This contributes to unavoidable excessive dispersion even in a purely Lagrangian transport approach, which

generates spurious transport across mixing barriers and steep tracer gradients. This results in the artificial generation of small scale structure and patchyness in the tracer fields. The mixing in the model has to be tuned to keep this patchyness at realistic levels. Hence, the mixing in the Lagrangian model will always tend to overestimate the real diffusion in the stratosphere somewhat, though to a much lesser degree than an Eulerian model. Eulerian models suffer from noise and discretization in the driving wind fields in a similar way, but additional numerical diffusion largely dominates these effects such that they do not become visible.

3 Validation and tuning

3.1 Introduction

The validation of ATLAS is based on the same aircraft data set that is used in the validation of CLaMS 3-D in Konopka et al. (2004) in order to facilitate direct comparisons: the model is validated by comparing a passive methane and a passive Halon-1211 tracer to values measured during the SOLVE/THESEO 2000 campaign in the Northern Hemisphere winter 1999/2000 (Newman et al., 2002). A number

of model runs is started with different mixing parameterizations. Each run simulates the period from 1 November 1999 to 15 March 2000. The methane and Halon tracers are initialized at 1 December. Between January and March, the model tracers are compared to in-situ measurement data from several high altitude flights conducted during the campaign that are able to temporally resolve fine tracer structures.

We will first describe the setup of the model runs and show some examples for model results. A quantitative validation will then be done by defining parameters for the agreement of the modeled and observed tracer-tracer relationships and the agreement of the spatial roughness of model and observations.

3.2 Setup

The model runs use the hybrid coordinate and are driven by meteorological data from the ECMWF ERA Interim reanalysis (Simmons et al., 2006, 2007) on 60 model levels (6 h temporal resolution, $2° \times 2°$ horizontal resolution). Heating rates (clear sky) are also taken from ERA Interim. ERA Interim data in the stratosphere has improved considerably compared to the ERA-40 reanalysis (Simmons et al., 2006, 2007), and, without any modifications, produces good agreement with observations in the following. Vertical diffusion coefficients are taken from the built-in climatology. The trajectory model uses a 4th order Runge-Kutta method for integration and a time step of 30 min. A change of the time step to 10 min does not lead to any significant improvements in the results.

Model runs are performed for all combinations of the resolutions r_0=50, 100, 150, 200 and 300 km and critical Lyapunov exponents λ_c=1, 2, 2.5, 3, 3.5, 4, 4.5, 5, and 10 d^{-1}. The number of air parcels ranges between 10^4 (300 km) and 10^7 (50 km). All runs at and below 150 km resolution are started both with the CLaMS and the ATLAS method of finding the next neighbors. For higher resolutions, only the CLaMS method can be used. Δt is set to 12 h and α to 250. Additionally, runs with fixed r_0 (100 km) and λ_c (3 d^{-1}) are started for different values of Δt or α. The effect of switching the random vertical coordinate on or off is also studied.

The lower model boundary is at the hybrid level 350 K and the upper boundary is around 1500 K. The number of virtual levels of the model (in the sense of the number of mixing depths Δz one needs to stack on top of each other to span the model domain) is 8 for 300 km resolution, 12 for 200 km, 16 for 150 km, 24 for 100 km, and 48 for 50 km.

The model is started on 1 November 1999, while the tracers are initialized on 1 December. This is done to avoid any spin-up effects in the mixing: usually, the equilibrium number of air parcels is lower than the number of air parcels the model is initialized with. This leads to a spin-up phase where more points are deleted in every model cycle than inserted, which in turn alters the mixing properties of the air parcels in an unrealistic way.

Outside the polar vortex, the methane tracer is initialized with monthly mean HALOE data from November and December 1999 (weighted equally) as a function of pressure and equivalent latitude φ_E (Grooß and Russell III, 2005). Inside the lower vortex ($\varphi_E > 64°$ N, $\theta < 712$ K), an OMS balloon profile from the LACE instrument measured at 19 November is used (Ray et al., 2002), since HALOE is not able to measure in high latitudes and darkness. The profile is corrected for the descent between 19 November and 1 December, as inferred from a model run driven by the heating rates described above. The profile above 712 K in the vortex is a profile from the highest equivalent latitude bin of HALOE scaled to match the LACE values below. The Halon-1211 tracer is initialized by a methane to Halon-1211 tracer-tracer relationship measured at the OMS balloon flight.

The model results for the methane tracer are compared to measurements of the ACATS-IV (Romashkin et al., 2001), ALIAS (Webster et al., 1994) and Argus (Loewenstein et al., 2002) instruments on 11 flights of the ER-2 aircraft between 20 January and 12 March (Newman et al., 2002). The Halon-1211 model results are compared to ACATS measurements on the same flights. Flights started from Kiruna (Sweden) and probed the stratospheric polar vortex and surf zone. Halon and methane measurements from ACATS are available every 70 s, corresponding to a resolution of about 20 km, while ALIAS and Argus measurements are available every 2 s (about 1 km resolution). Wherever possible, we use the high resolution data, tracer-tracer relationships and Halon plots use the low resolution.

3.3 Examples

Figure 6 shows model results and measurements for the 27 January flight as an example. Results are shown for 50 km resolution and optimal mixing parameters (λ_c=3 d^{-1}, α=250, Δt=12 h, see Sect. 3.5). Only the highest tested resolution of 50 km is able resolve the fine filamentary structures at the edge of the polar vortex. The flight crosses the vortex boundary twice at a mean flight level of 50–60 hPa and contains a dive to 100 hPa around 12:00 UTC. For the comparison, the flight path is transformed to the position of the probed air parcels at 12:00 UTC by calculating short forward or backward trajectories starting at every measurement. Model results at 12:00 UTC are interpolated to the transformed measurement positions by averaging the mixing ratios over the next model neighbors of every measurement location, with a 2-D triangulation performed in the same way as in the mixing step of the model. The model values in the average are weighted according to their distance to the measurement: the weight is 1 for distances smaller than $2r_0$ and decreases linearly to zero at $3.5r_0$.

Panels a and b show measurements and model results for methane and Halon as a function of time, while Panels c and d show the same as a function of pressure. The agreement of the modeled and observed methane and Halon-1211 values is
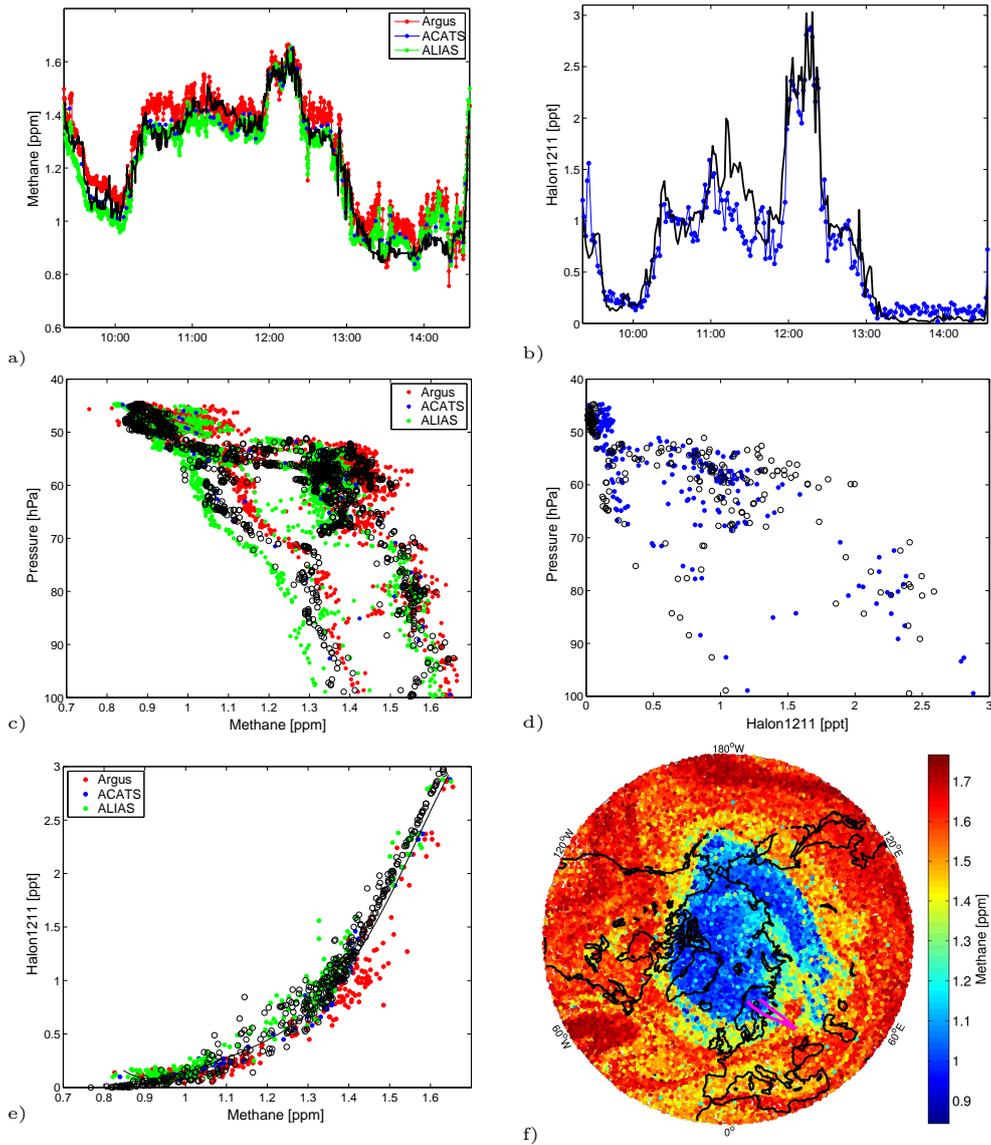
Fig. 6. Example comparison of model results (r_0=50 km, λ_c=3 d^{-1}) and flight data for 27 January, 12:00 UTC: **(a)** Measured methane mixing ratios (red: Argus, blue: ACATS, green: ALIAS) and modeled methane mixing ratios (black line) as a function of flight time (UTC), **(b)** Measured Halon-1211 mixing ratios (blue: ACATS) and modeled Halon-1211 mixing ratios (black line) as a function of flight time, **(c)** Methane mixing ratio as a function of pressure (colored dots: measurements, black circles: model values), **(d)** Halon-1211 mixing ratio as a function of pressure, **(e)** Tracer-tracer relationship of methane and Halon-1211, **(f)** Modeled tracer field for the Northern Hemisphere for all air parcels (dots) between 432–444 K. The magenta line denotes the positions of the air parcels probed by the flight at 12:00 UTC. All model values have been interpolated from the nearest model parcels to the transformed flight path.

quite good, both for the gradient across the vortex edge and the vertical structure. Inner and outer vortex air masses can clearly be distinguished both in the model and the observations below 70 hPa and the gradual transition from inner to outer vortex values above 70 hPa is also matched quite well by the model. The modeling of Halon is challenging, since the mixing ratios of Halon very sharply drop to zero above 50 hPa. Hence, the location of the drop is a good test for the accuracy of the heating rates, the accuracy of the winds and the mixing parameterization. Panel e shows the modeled and

observed tracer-tracer relationships, which we will use later to validate the model. Panel f shows the methane tracer field at the flight level, with the polar vortex characterized by low methane values.

The interpolation to the probed air parcels of the flight will inevitably lead to additional diffusion in the model results, and different approaches for finding corresponding model values show considerable differences in the results, see also Konopka et al. (2003). Figure 7 shows the remarkable differences for the flight on 11 March, which crossed a filament
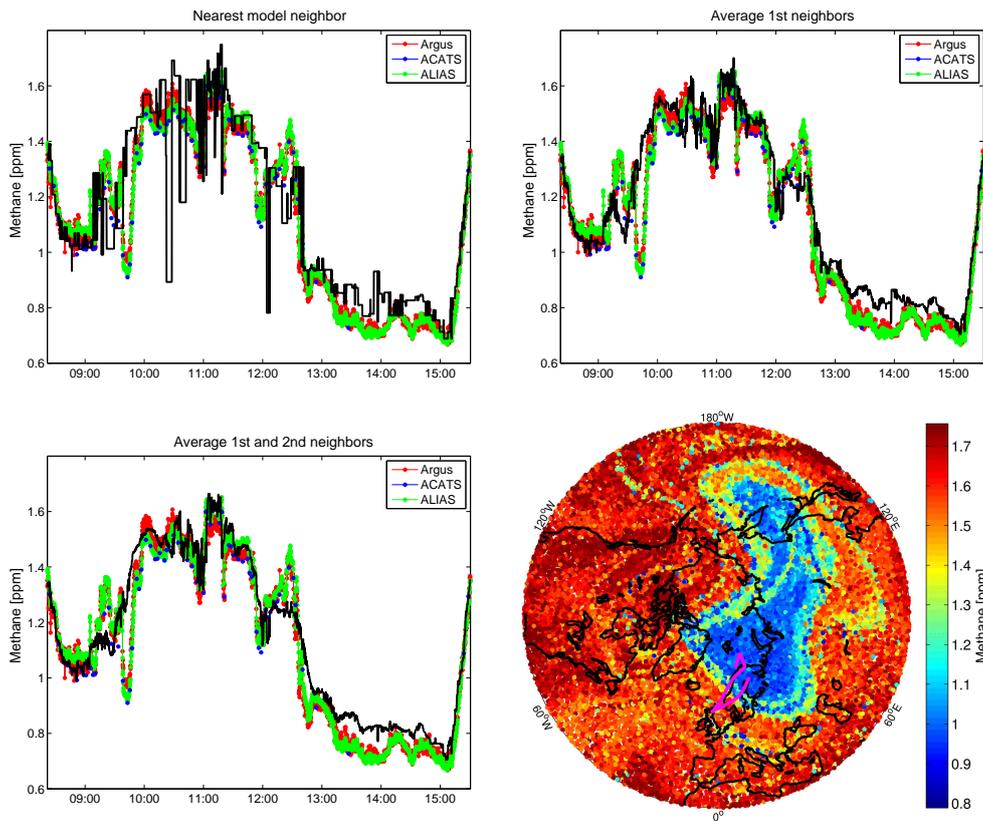
Fig. 7. Different methods of the interpolation to the flight path. Measured methane mixing ratios (colored dots) and modeled methane mixing ratios (black line) as a function of flight time (UTC) for the 11 March flight (λ_c=3 d^{-1}, r_0=50 km): for the nearest model neighbor of the measurement (upper left), the average of the first neighbors (upper right) and the average of the first and second neighbors (lower left). More averaging means a better agreement in areas of low gradients (due to reduced scatter), while less averaging better preserves fine filament structures. Lower right: modeled tracer field for the Northern Hemisphere for all air parcels (dots) between 432–444 K.

of extra-vortex air inside the vortex twice during the flight. While the filament is clearly visible if using only the nearest model neighbor to each measurement as the corresponding model value (free of any interpolation), it is more difficult to see for the average over the next neighbors and barely visible for the average over the next and second neighbors. On the other hand, parts of the flight where gradients are low show a better agreement with more averaging in the interpolation. This effect gets more prominent if λ_c is further increased (not shown). In this sense, the interpolation must be considered as an integral part of the modeling. Hence, for the actual validation, we try to find measures free of any additional interpolation.

Figure 8 demonstrates how the results change if the mixing parameter λ_c is varied. The left column shows an example where the mixing in the model is stronger than observed, leading to a mixing curve in the model that is on the concave side of the observed mixing curve. The middle panel shows an example for optimal mixing parameters and the right panel shows an example for too low mixing, where the modeled mixing curve is on the convex side of the observed curve.

3.4 Validation measures

The quality of the model runs is judged by calculating a parameter ε for the agreement of the modeled and the observed tracer-tracer relationships for every run. Likewise, a parameter γ for the agreement of the spatial "roughness" of the observed and modeled tracer fields is calculated for every run.

While the parameter ε is a straightforward measure for the mixing intensity of the model, the parameter γ needs some explanation. For low mixing intensities (high λ_c), the model fields show scatter and filaments, where the measurements do not show any. There are two possible reasons for this: first, tracer structures that would be mixed to background values in reality are not mixed in the model due to the low mixing intensity. But additionally, the spurious transport described in Sect. 2.3.6 generates small scale structure and leads to the same effect, see also Legras et al. (2005). The two parameters ε and γ give some possibility to distinguish between the two effects: spurious transport generates small scale structure and hence has an effect on γ. But these transport effects do not result in changes in the tracer relationships that are the
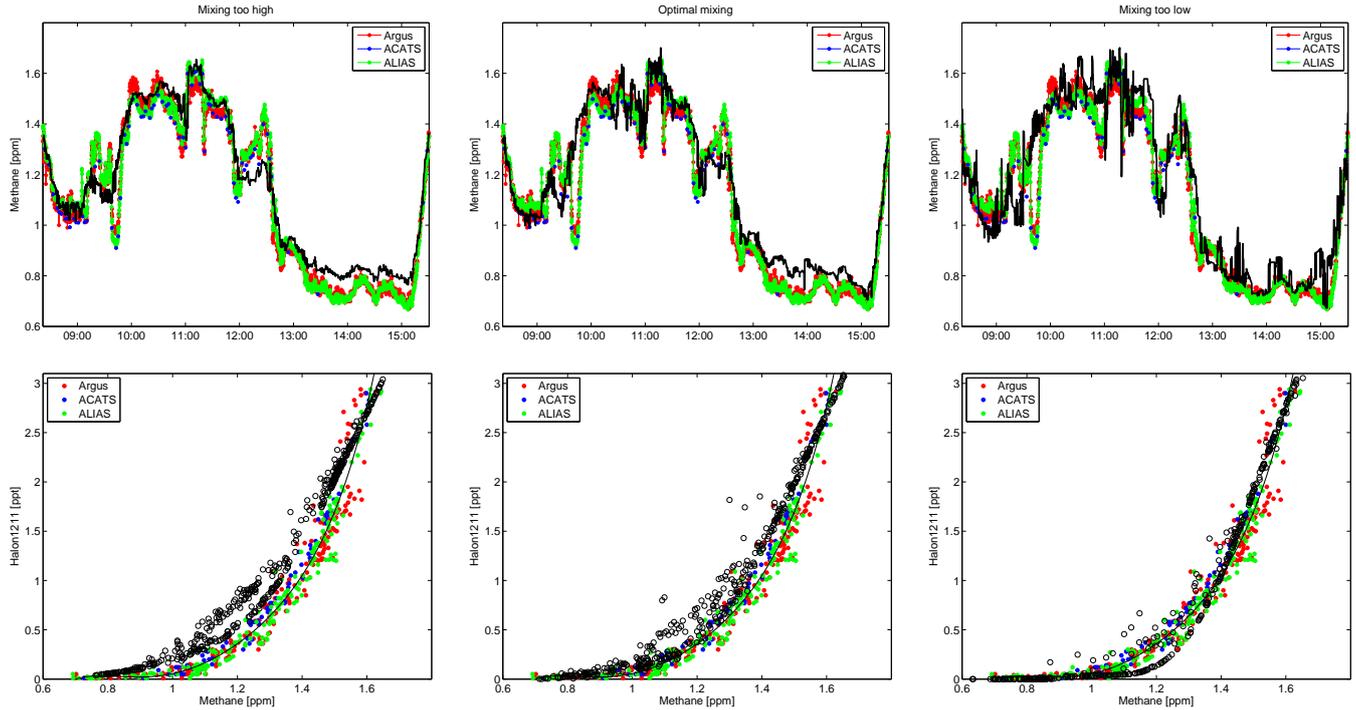
Fig. 8. Change of results when varying λ_c. Top: measured methane mixing ratios (colored dots) and modeled methane mixing ratios (black line) as a function of flight time (UTC) for 11 March 2000, r_0=50 km and λ_c=2 d^{-1} (left, too much mixing), λ_c=3 d^{-1} (middle, best agreement with observations) and λ_c=5 d^{-1} (right, almost no mixing). Bottom: same for the tracer-tracer relationship.

basis for ε in first approximation, since the artificial filaments will be mixed with the same intensity as any filament. Any exaggeration of model mixing used to balance the generation of patchyness and to tune γ to values close to observations results in deviations of modeled and observed tracer relationships which show up in ε.

In the following, we will denote the ith methane measurement with m_i (counting over all flights without resetting the counter), and the corresponding Halon-1211 measurement with h_i.

The parameter ε for the comparison of the mixing intensity is defined by the sum of the distances in tracer-tracer space of the measurement pairs $(m_i^{\mathrm{observed}}, h_i^{\mathrm{observed}})$ to the corresponding model pairs $(m_i^{\mathrm{model}}, h_i^{\mathrm{model}})$ in Konopka et al. (2004). That is,

$$\varepsilon = \sum_{i=1}^{N} (m_i^{\mathrm{observed}} - m_i^{\mathrm{model}})^2 + (h_i^{\mathrm{observed}} - h_i^{\mathrm{model}})^2.$$

A sketch of the method is shown in Fig. 9 (left). Since the absolute magnitude of the sum for methane and the sum for Halon-1211 is very different, relative differences are taken in Konopka et al. (2004)

$$\varepsilon = \sum_{i=1}^{N} \left(1 - \frac{m_i^{\mathrm{model}}}{m_i^{\mathrm{observed}}}\right)^2 + \left(1 - \frac{h_i^{\mathrm{model}}}{h_i^{\mathrm{observed}}}\right)^2 \qquad (11)$$

where N is the total number of measurements over all flights. Small values mean a good agreement between the mixing

intensity of the model and the observations. Corresponding model pairs $(m_i^{\mathrm{model}}, h_i^{\mathrm{model}})$ are found by looking for the nearest model neighbor to the measurement location in Konopka et al. (2004) to avoid interpolation to the measurement location. Interpolation would cause additional mixing and may considerably alter the tracer-tracer correlation curve (e.g. by introducing spurious mixing lines, not shown).

There are some possibilities to improve this definition: first, for mixing ratios near zero (as it is the case for Halon-1211 in several parts of the flights), the relative differences go to infinity, and some single measurements can erroneously dominate the overall sum due to measurement error.

Secondly, errors in the wind fields used to drive the model can cause the position of some filament in the model to be shifted from its position in the observations to a slightly different place. This would cause a large distance of the corresponding (m_i, h_i) pairs in tracer-tracer space (where one pair is in the filament and the other is not). However, the mixing intensity could still be right, that is the modeled mixing ratios in the filament would match the observed tracer mixing ratios and the filament would have the right extent and shape. Hence, in a model with perfect mixing but slight systematic errors in the driving wind fields, the mixing curve of the model would still look exactly like the observed mixing curve. Since ε should only measure the mismatch in mixing intensity and not displacements of the filaments, this is not optimal. Konopka et al. (2004) argues that for too low
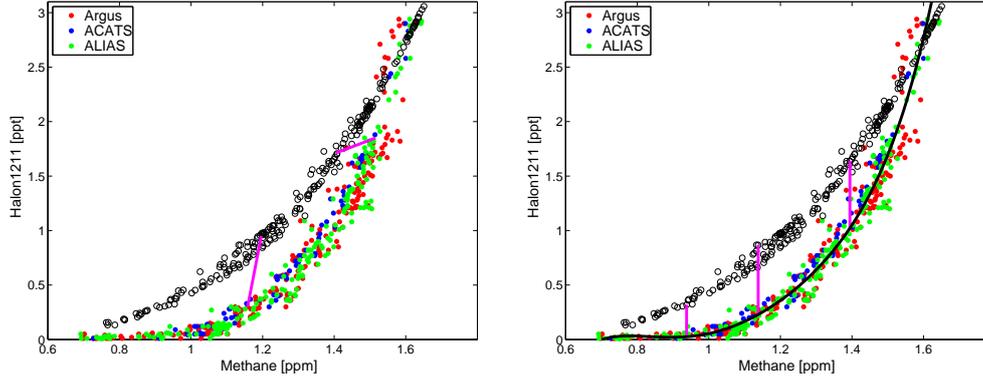
Fig. 9. Sketch of the two methods for calculating the parameter ε. Left: in the definition of Konopka et al. (2004), ε is the sum of the distances of corresponding individual points in the Halon-methane tracer-tracer space. Right: in this work ε is the sum of the distances of modeled values of Halon-1211 to a curve fitted through the observed Halon-1211 values as a function of observed methane.

mixing intensities, spurious filaments that would be mixed in the real atmosphere would also cause large distances of corresponding points in tracer-tracer space, which outweigh the effect of displaced filaments in ε. We argue that the key measure of mixing properties of the model is the shape of the tracer-tracer relation and not the point-to-point distances between model and observations in tracer-tracer space and use an approach that results in a clearer separation of the roughness effect (which is not only caused by too low mixing, but also by the random errors in wind fields) and the effect of displaced filaments from the mismatch in mixing intensity in tracer-tracer space. Hence, we try to put the roughness effect only into γ and try to define ε such that it is only affected by the mixing intensity.

A suitable definition minimizes if the shape of the mixing curve of the model gets more and more similar to the shape of the observed mixing curve. We also look for a definition that allows us to take the model values directly to avoid interpolation. This can be accomplished by the absolute differences of the modeled Halon-1211 values to a curve fitted through the observed Halon-1211 as a function of observed methane. A sketch of the method is shown in Fig. 9 (right). The sum of these differences is largely proportional to the area between the mixing curves of the model and the observations. The observed methane is taken from the average of the three instruments. To obtain the absolute difference, we evaluate the fit function at the modeled methane value. That is

$$\varepsilon = \frac{1}{M} \sum_{j=1}^{M} |h_j^{\text{model}} - \hat{h}_{k(j)}^{\text{obs}}(m_j^{\text{model}})| \tag{12}$$

where $k(j)$ is the flight number belonging to the modeled value j and $\hat{h}_k^{\text{obs}}$ is a polynomial fit through the observed tracer-tracer relationship of flight k. Since we use a continuous curve for the observed tracer-tracer relationship, we do not need to use a single corresponding value m_i^{model} for every measurement m_i^{observed}. Instead, we use all model air parcels

that are next neighbors of the measurement locations, from the 2-D triangulation used in the mixing step and with Δz centered at the measurement location. This method gives us a cloud of air parcels along the flight path. This is indicated here by using j for i as the index and M for N for the number of points. A potential disadvantage of the method is that it can only be applied to data only containing one compact mixing curve (and no additional mixing lines).

In some cases individual methane values in the model are outside of the range of the measured methane values, where the fitted curve is valid. In this case, we do not use these values, so that ε is not based on exactly the same set of points for some model runs. The validity of this approach is verified by extending the mixing curves with a best guess of the likely missing methane values and calculating the parameter again (these tests were not used for the validation). Results are almost identical. The best guess is simply obtained by extending the data of the tracer curves with methane and Halon-1211 data from the third flight for the 5 other flights at the end of January or start of February, since it has the largest range of methane values. The flights at the end of February and in March are extended with data from the eighth flight.

The parameter γ for the roughness of the field in Konopka et al. (2004) is defined in two steps: the sum of the absolute differences of successive methane measurements is

$$\Delta_{\text{observed}} = \sum_{i=1}^{N-1} |m_{i+1}^{\text{observed}} - m_i^{\text{observed}}| \tag{13}$$

The same can be done for the model values of methane. The ratio of the measured and modeled Δ is

$$\gamma = \frac{\Delta_{\text{model}}}{\Delta_{\text{observed}}} \tag{14}$$

Values near 1 are defined as the optimum. With this definition the magnitude of the roughness will depend on the spatial scale of the structures we are looking at. It will make

a difference if the measurement locations are 1 km or 20 km apart. More importantly, it will also make a difference if the model air parcels are 50 km or 300 km apart. In this sense, γ is not a very robust measure of the roughness. Hence, we use a definition of γ that compares the roughness $\Delta_{\mathrm{observed}}$ and Δ_{model} at the spatial scale of the model resolution and is robust to changes in the spatial distance of the measurements.

Our approach compares the standard deviations of the modeled and observed methane values at the spatial scale of the effective model resolution: we determine the next model neighbors of each measurement i (exactly as described for the other parameter ε) and calculate the standard deviation σ_i of the modeled methane values for each set of next neighbors (avoiding any interpolation). Now, all standard deviations are averaged to give a mean standard deviation σ_{model}. In addition, we calculate the mean distance of the points in every set of next neighbors and compute its average $\bar{d}$. Now, we look for all pairs $(m_i^{\mathrm{observed1}}, m_i^{\mathrm{observed2}})$ of measurements which are separated by a distance within $[\bar{d}-\Delta d, \bar{d}+\Delta d]$, with $\Delta d = 5$ km. It can be shown that the standard deviation of this set of points is given by the following equation:

$$\sigma_{\mathrm{observed}} = \sqrt{\frac{1}{2N} \sum_i^N (m_i^{\mathrm{observed1}} - m_i^{\mathrm{observed2}})^2} \qquad (15)$$

where N is the number of pairs. Now, we define

$$\gamma = \frac{\sigma_{\mathrm{model}}}{\sigma_{\mathrm{observed}}} \qquad (16)$$

in equivalence to Eq. (14). Theoretically, a term for the statistical measurement error in the observed values needs to be considered, since the standard deviation of the observations will not only be caused by real variations of the mixing ratio but also by measurement noise. This would cause an offset σ_{bias} to $\sigma_{\mathrm{observed}}$ and the denominator would be $\sqrt{\sigma_{\mathrm{observed}}^2 - \sigma_{\mathrm{bias}}^2}$. However, it turns out that this correction is negligible.

It is important to note that our γ gives us only a measure for the deviation of the model variability from what we would expect from the observations at the given model resolution. There is no penalty if small filaments just do not appear in coarser model resolutions. In this sense it gives a relative but not absolute measure for the expected variability.

Note that our definition of ε and γ does not imply a perfect agreement of the observed and modeled time series along the flight path if $\varepsilon = 0$ and $\gamma = 1$. For example, all filaments could be displaced by a small distance between the modeled and observed time series along the flight path in this case.

3.5 Validation results

Figure 10 shows the results for the parameters γ and ε for different model resolutions r_0 and different values of the critical Lyapunov exponent λ_c (for a comparison of our definition of ε and γ and the one given in Konopka et al., 2004, see Appendix B).

For a given r_0, data points show more intense mixing to the left (low values of λ_c) and less intense mixing to the right (high values of λ_c). $\lambda_c = 10$ d^{-1} is virtually identical with no mixing at all ($\lambda_c \to \infty$).

The agreement of the modeled tracer-tracer relationship with the observed one generally improves with increasing λ_c. For higher resolutions it does only reach a shallow minimum at roughly $\lambda_c = 3$ d^{-1} to $\lambda_c = 4$ d^{-1} (see also Fig. 12), which compares well with the value $\lambda_c = 3$ d^{-1} given in Konopka et al. (2004). This basically means that the observed mixing is extremely low and very anisotropic: the ratio between r_+ and r_- is larger than 20 for these values of λ_c. Switching mixing off in the model is actually not a bad option for optimizing ε. Indeed, the observed tracer-tracer relationships on the first flight on 20 January and the last flight on 12 March are very similar, indicating that mixing in the real stratosphere is extremely low.

Generally, the agreement of the tracer-tracer relationships improves with higher resolution r_0 if λ_c is held constant. This is due to the fact that the model gets more diffusive with increasing distance between the mixed air parcels.

Where the triangulation approach of the ATLAS model is applicable (dashed lines), it is generally less diffusive than the CLaMS approach (solid lines), see also Fig. 14. For the same set of mixing parameters, the results of the ATLAS triangulation approach are superior to the CLaMS approach. One likely reason for that is that the number of mixed parcels is lower in the ATLAS approach than in the CLaMS approach for the same λ_c (Fig. 11). This is due to the lower average vertical distance between mixed air parcels: the maximum distance of two mixed air parcels is Δz in the CLaMS approach (if both are at the edges of the layer), but only $\Delta z/2$ in the ATLAS approach (since one of the parcels is by definition at the center of the local layer). A possibility to work against that in the CLaMS approach would be to increase λ_c until the percentage of mixed parcels is low enough. However, there would still be qualitative differences between the approaches. E.g., if an air parcel is situated at the boundary of a layer in the CLaMS approach, it is possible that it does not mix with a vertically close parcel just because it is in the adjacent layer. In this respect, the ATLAS approach behaves qualitatively more like in the real atmosphere. In fact, there are some runs with the ATLAS approach for 150 km resolution, which show a lower ε than any run with the CLaMS approach in the same resolution, no matter what λ_c is used.

However, note that for coarse resolutions (say 300 km) the mixing process in the ATLAS as well as the CLaMS approach will become unrealistic. The diffusion coefficients of single mixing events will become unrealistically large and that can only be compensated by setting λ_c to large values to obtain realistic bulk diffusivities (see Sect. 4). But this will still be a better option for long-term runs in coarse resolutions than using an Eulerian approach.
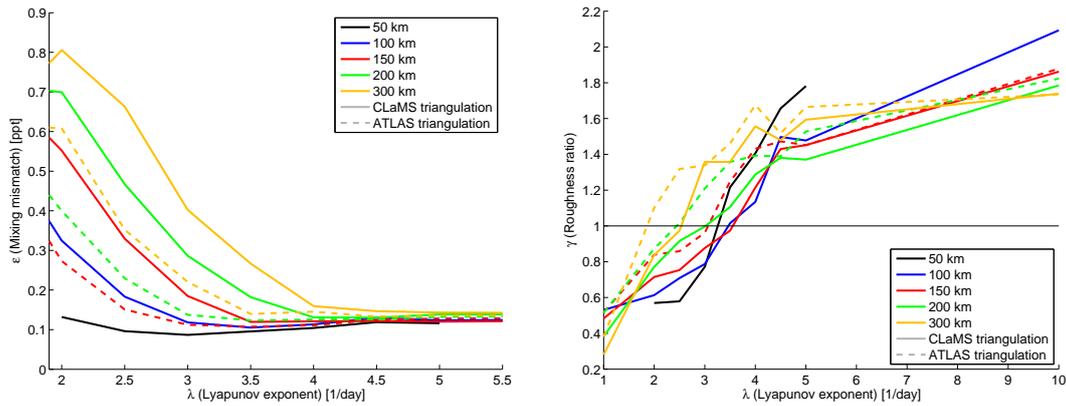
Fig. 10. Mixing mismatch and tracer roughness. Left: the parameter ε for the difference of the modeled and observed tracer-tracer relationship for different values of the critical Lyapunov exponent λ_c (axis) and resolution r_0 (colors). Solid lines show results for the triangulation approach of CLaMS, dashed lines show results for the triangulation approach of ATLAS. Right: same for the parameter γ for the tracer roughness ratio.
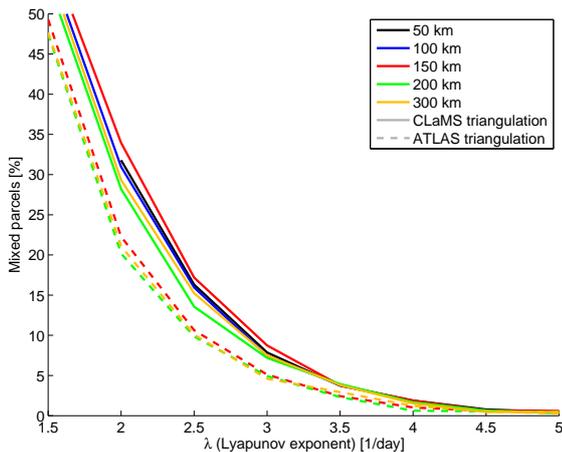

Fig. 11. Percentage of mixed parcels for different values of the critical Lyapunov exponent λ_c (axis) and resolution r_0 (colors).
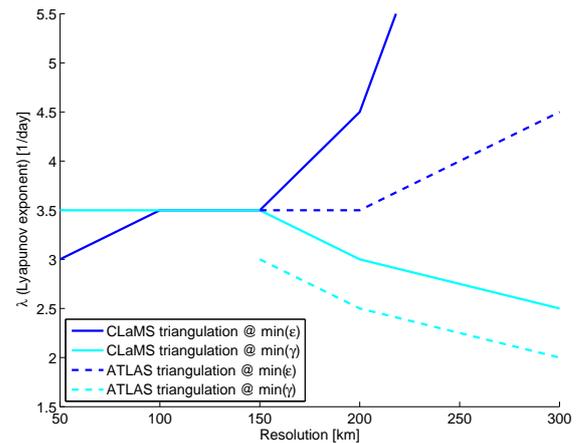

Fig. 12. Critical Lyapunov exponents where the mixing mismatch ε minimizes (blue) and the roughness ratio γ is closest to one (cyan). Solid lines show the CLaMS triangulation approach, dashed lines show the ATLAS triangulation approach.

For all resolutions, the roughness ratio γ increases more or less uniformly with decreasing λ_c (the model fields are too smooth for low values of λ_c and too rough for high values of λ_c, see Fig. 10 (right). The optimal value for λ_c based only on $\gamma \approx 1$ is about $2\,d^{-1}$ to $3.5\,d^{-1}$ for all resolutions and both triangulation approaches.

Since the effect of random errors in the wind fields and numerical discretization shows up mainly in the roughness ratio γ and not in the mixing mismatch ε (because the artificial filaments would still be mixed with the right intensity), γ should optimize at lower values of λ_c (at higher mixing intensity) than ε. Figure 12 shows that this is in fact what is observed in most validation runs. It shows the λ_c values where the ε curves minimize (blue) and γ is closest to one

(cyan). Note however that the minima are shallow and that some uncertainty is introduced by this. In addition, the ATLAS triangulation approach is optimal at smaller values of λ_c than the CLaMS approach, possibly since it is less diffusive and needs more mixing in the model to obtain optimal results. The values at $300\,km$ should not be taken too seriously, since the minimum of ε is at the highest tested λ_c of $10\,d^{-1}$ here. Often, it will be unavoidable to find a compromise in optimizing γ and ε, i.e. to fight the patchyness by setting the mixing intensity to slightly higher values than suggested by ε (lower values for λ_c).

The results are not very sensitive to the value of the aspect ratio α. Values between 200 and 500 give comparable results to 250. As long as the product $\lambda_c \Delta t$ stays constant, the re-

sults are also not very sensitive to the mixing time step Δt. Values of 6 h and 24 h give comparable results to 12 h. These results are also confirmed by the validation of the CLaMS model (Konopka et al., 2004, 2005).

Figure 13 shows that the effect of adding the random vertical coordinate is small in the range of the minimum ε values both for the ATLAS and CLaMS triangulation. However, switching the random vertical coordinate off does decrease the diffusivity for the lowest λ_c values in the ATLAS triangulation approach. This is very likely due to the vertical clustering into layers of enhanced parcel density that sets in if the random coordinate is switched off and the mixing intensity is high. This effectively suppresses vertical mixing and lowers the effective diffusion coefficient to values which are normally observed at higher λ_c values. That is, while improving the vertical homogeneity of the model, the random coordinate has either negligible effects on the diffusivity (high λ_c) or is just necessary for a realistic behaviour of vertical mixing in the ATLAS approach (low λ_c).

Fig. 13. Effect of adding a random number to the vertical coordinate of newly inserted parcels. The solid blue line shows ε for 100 km resolution, the CLaMS triangulation and the random vertical coordinate on, the dashed red line for 150 km, the ATLAS triangulation and the random vertical cordinate on. Dotted lines are the same with the random vertical coordinate off.

4 Estimation of diffusion coefficients

It is desirable that the diffusion coefficients obtained with the optimal mixing parameters are in the range of other estimates. However, there are few direct and independent estimates of diffusion coefficients (e.g. from radar measurements, like Fukao et al., 1994, see also Wilson, 2004). Most estimates (including ours) are based on indirect methods as the reconstruction of observed tracer filaments, profiles or tracer-tracer relationships, usually with a diffusion component and sometimes a Langrangian transport component, e.g. by a simple 1-D diffusion model for the mean vertical profile (Massie and Hunten, 1981), numerical models for strain and diffusion of filaments (Balluch and Haynes, 1997; Waugh et al., 1997), or by averaging stochastic back trajectories (Legras et al., 2005). In contrast to e.g. Balluch and Haynes (1997), the information about the diffusion coefficient comes mainly from the tracer-tracer relationship in our method and not from reconstructing fine filaments, which are only resolved in the highest resolutions. A general problem of a comparison are the large differences in the methods. Some are local (radar on the 100 m scale), some give bulk estimates, some are Eulerian, some Lagrangian, some direct and some indirect.

Figure 14 shows the effective mean (bulk) vertical and horizontal diffusion coefficients north of 60° N at 430 K as a function of λ_c and r_0, calculated as in Sect. 2.3.5. Vertical diffusion coefficients obtained from the studies mentioned above are also shown for comparison. The modeled vertical diffusion coefficient increases with lower λ_c, which is due to the increasing number of mixing events, and increases with coarser resolution, which is due to increasing distances between mixed parcels. Figure 15 shows the value of the vertical diffusion coefficient at the optimal values of ε
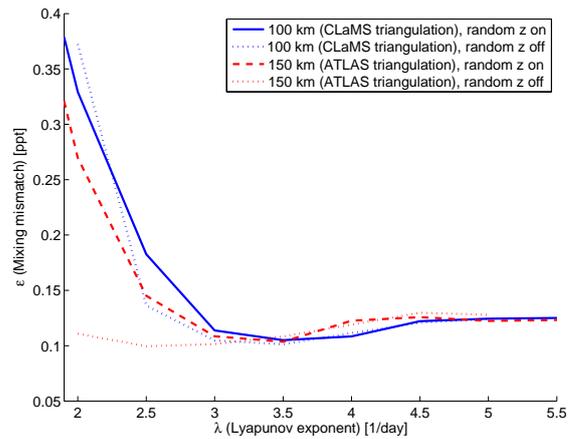
and γ in the same manner as in Fig. 12. The vertical diffusion coefficient takes values between 0.07 m²/s and 0.36 m²/s for the optimal ε and the CLaMS triangulation approach, and values between 0.02 m²/s and 0.1 m²/s for the optimal ε and the ATLAS triangulation approach. The optimal γ shows vertical diffusion coefficients that are probably too large, if compared with the estimates from the other studies. These values are increased due to errors in the wind fields. In summary, our results suggest a value of about 0.1 m²/s for the vertical diffusivity. As noted, the estimated bulk diffusivities will probably be more realistic for high resolutions, since the diffusivities of single mixing events are probably too large for coarser resolutions.

Our estimate agrees well with the value 0.1 m²/s given for the lower stratospheric surf zone in Legras et al. (2005). Values from the radar measurements (0.1 m²/s to 0.5 m²/s for Fukao et al., 1994) tend to be larger than other estimates. Also the Massie and Hunten (1981) value of 0.58 m²/s at 22 km seems to be relatively large. As estimated in Sect. 2.3.3, values like these would roughly correspond to a minimum width of a filament of 50 km before it disappears. One reason might be that the effects of the large-scale advection by the Brewer-Dobson circulation are not considered in studies like Massie and Hunten (1981). On the other hand, the estimate of Balluch and Haynes (1997) (0.01 m²/s to 0.001 m²/s) and of Legras et al. (2005) for the inner vortex (0.01 m²/s) are on the lower side of our estimate. Differences can at least partly be explained by the likely dependence on season and location (e.g. single filaments in Balluch and Haynes, 1997, or the average north of 60° N in our approach compared to the separate values for surf zone and vortex in Legras et al., 2005).
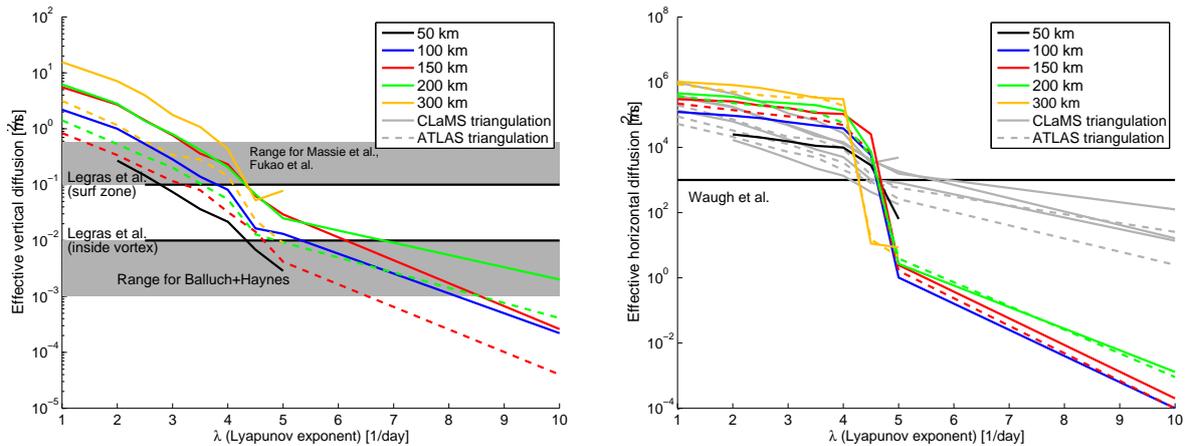
Fig. 14. Left: effective vertical diffusion coefficient north of $60°$ N at a layer around 430 K for different values of the critical Lyapunov exponent λ_c (axis) and resolution r_0 (colors). Solid lines show results for the triangulation approach of CLaMS, dashed lines show results for the triangulation approach of ATLAS. Shaded areas and solid black lines show vertical diffusion coefficients from other studies for comparison: Massie and Hunten (1981), Fukao et al. (1994), Balluch and Haynes (1997), Legras et al. (2005). Values for $\lambda_c=10\,d^{-1}$ and $r_0=300\,km$ are zero and not shown. Right: same for the effective horizontal diffusion coefficient. Grey lines are the vertical diffusion coefficients from the left panel scaled by α^2. Solid black line shows estimate from Waugh et al. (1997).
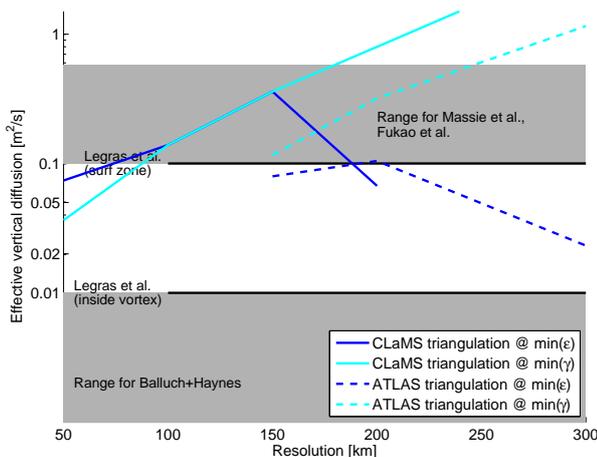

Fig. 15. Value of the vertical diffusion coefficient where the mixing mismatch ε minimizes (blue) and the roughness ratio γ is closest to one (cyan). Solid lines show the CLaMS triangulation approach, dashed lines the ATLAS triangulation approach. Shaded areas and solid black lines show vertical diffusion coefficients from other studies for comparison as in Fig. 14.

The right panel of Fig. 14 shows the horizontal diffusion coefficient, which ideally should be a scaled version of K_z, since Haynes and Anglade (1997) claims $K_h=\alpha^2 K_z$. The grey lines show the vertical diffusion coefficients from the left panel scaled by α^2. In general, the agreement is quite good for values smaller than $\lambda_c=5\,d^{-1}$, although the horizontal diffusion coefficients are somewhat larger than expected from K_z (note the logarithmic scale). This could be interpreted as the effective aspect ratio of the model being somewhat larger than 250. The modeled horizontal diffusion coefficient also increases with lower λ_c, which is now

not only due to the increasing number of mixing events, but also modified by the increasing diffusivity orthogonal to the flow direction and the decreasing diffusivity in flow direction (Sect. 2.3.5). Waugh et al. (1997) gives a value of $10^3\,m^2/s$ for the horizontal diffusion coefficient, which seems to be somewhat lower than our effective diffusion coefficients. The sharp drop in the horizontal diffusion at λ_c values of $5\,d^{-1}$ and $10\,d^{-1}$ is due to the fact that only very few mixing events occured at these values and that all events were merging events (r_-), while at lower λ_c values, also r_+ events occured, which have a much larger diffusion coefficient (at the same mixing parameters). That hints at the fact that the mixing in the model will get unrealistic at these large values of λ_c.

5 Conclusions

The new global Chemical Transport Model (CTM) ATLAS with Lagrangian transport and mixing was presented. In this paper, we focussed on a basic model description and the presentation and validation of the transport and mixing module of the model. In conjunction with a projected second part concentrating on the chemistry module, the paper is thought to serve as the reference citation for the model.

The ability of the transport and mixing model to reproduce observed tracer data and fine-scale tracer structure has been successfully demonstrated. Excellent agreement to observed structures up to scales of typical polar vortex filaments can be reached with suitable settings of the free parameters of the model, which were extensively tested and validated.

The results of the validation suggest a vertical diffusion coefficient on the order of $0.1\,m^2/s$ in the high-latitude lower stratosphere, which fits well into other estimates. However,

the vertical diffusion coefficient remains a parameter that can only be estimated with large uncertainty, ranging from $0.5\,\mathrm{m^2/s}$ to $0.001\,\mathrm{m^2/s}$ in the literature cited here.

The model makes heavy use of concepts developed for the CLaMS model, and consequently part of the study focussed on comparison with that model. Generally, results seem to be of comparable quality and the changes introduced in AT-LAS were demonstrated to improve the results (new concept for identifying the neighbors for the mixing process) or to have little effect (effect of random vertical coordinate in the CLaMS triangulation, diffusion coefficient climatology versus entropy-based diffusion coefficient).

As always, time and space did not permit to study all aspects of the model and the validation. For example, since the diffusivity of the atmosphere will vary in time and space, it would be interesting to validate the model with data from other geographical regions or seasons, e.g. to see if the model is also able to reproduce mixing lines in cases of strong mixing events (i.e. in the case of fast anomalous mixing, compared to the constant mixing which is slowly lifting the correlation curve on the concave side). It would also be illuminative to compare the validation results of the Lagrangian model to the results an Eulerian transport code would give.

Future plans with the model include the improvement of the representation of the troposphere, and in the long run, the full coupling to a Atmosphere Ocean General Circulation Model.

Appendix A

Some details of the ATLAS triangulation method

The 3-D triangulation is performed in cartesian space, that is the spherical coordinates are transformed to cartesian coordinates before triangulation with a given radius of the Earth of $a=6371\,\mathrm{km}$. Since the large aspect ratio α of the shallow atmosphere would produce very flat tetrahedrons and severe numerical problems in the triangulation, the atmosphere is scaled by a factor of 250 before triangulation. That is

$$
\begin{aligned}
x &= (a+250z')\cos\varphi\cos\lambda \\
y &= (a+250z')\cos\varphi\sin\lambda \qquad\qquad\qquad \text{(A1)}\\
z &= (a+250z')\sin\varphi
\end{aligned}
$$

where z' is log-pressure altitude, λ longitude, φ latitude, x, y, z cartesian coordinates. Finally, an artificial point is inserted at the center of the Earth to avoid that the triangulation creates tetrahedrons in the Earth's interior. This point is later removed from all neighbor relationships. A visualization of the method is shown in Fig. A1.

The 3-D triangulation is only needed to produce a cloud of points surrounding the original point, which includes all points that could possibly be a neighbor of the original point
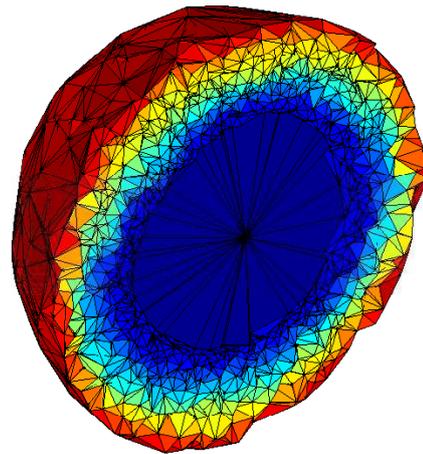

Fig. A1. 3-D triangulation of air parcels on the globe in the ATLAS triangulation method. Edges connecting air parcels are shown as black lines, colors denote an arbitrary altitude scale for better visualization. The sphere is cut into two halfes, with one half removed, to allow for a look into the interior of the triangulation.

in the final 2-D triangulation. Empirically it is found that including the first and second neighbors makes the cloud large enough for this purpose.

Appendix B

Comparison of the different definitions of ε and γ

Figure B1 shows ε calculated by Eq. (11) (originally from Konopka et al., 2004) for comparison with our definition of the parameter by Eq. (12) in Fig. 10 (left). ε according to Eq. (11) (with corresponding model points obtained by nearest neighbors) shows more scatter than our ε. In addition, some of the lines for different resolutions are cutting across each other (note that Fig. 11 in Konopka et al. (2004) uses a smoothed version of the curves). However, it is reassuring that the general conclusions are still the same if using Eq. (11).

Figure B2 shows the ATLAS γ (Eq. 16) for $r_0=100\,\mathrm{km}$ as reference, and γ as defined by Eq. (14) based on three different methods for finding the corresponding model point m_i^{model} to the measurement: nearest model neighbor of the measurement and two methods of interpolation, namely average over the next neighbors and average over the next and second neighbors. While the method without averaging (nearest neighbor) shows too much variability for all λ_c and never reaches the optimal value 1, the method with the most averaging (first and second neighbors) shows relatively small values at large values of λ_c, since it tends to smear out filaments that are actually in the model data and introduces spurious mixing. The blue lines show the results if the low resolution measurements of ACATS are used (about $20\,\mathrm{km}$
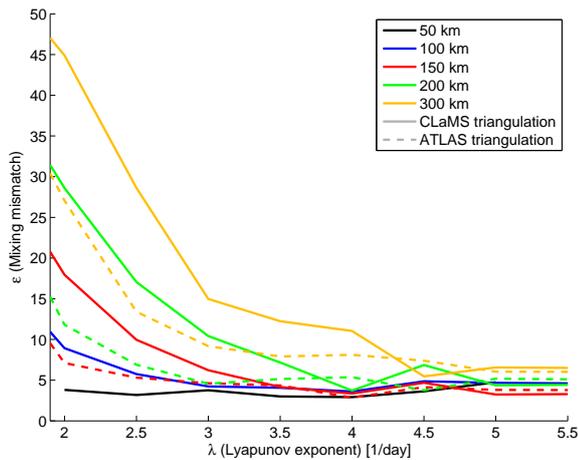
Fig. B1. The parameter ε as defined in Eq. (11) (originally from Konopka et al., 2004) with nearest neighbors as corresponding model points, for comparison with ε as defined in Eq. (12) shown in Fig. 10 (left).
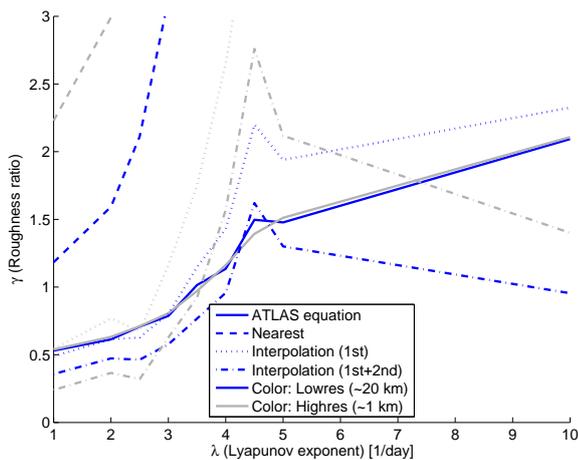

Fig. B2. The parameter γ only for r_0=100 km. The solid lines show the definition for γ from this study (Eq. 16), the different dashed lines show the definition of γ in Konopka et al. (2004) (Eq. 14) with different approaches to obtain the model points m_i^{model}: nearest neighbor of measurement, average over first neighbors of the measurement, average over first and second neighbors. The blue lines show results if only using the ACATS measurements (about 20 km resolution) and the grey lines if only using the Argus and ALIAS measurements (about 1 km resolution).

resolution) and the grey lines show the results if the high resolution measurements of Argus and ALIAS are used (about 1 km resolution). While the γ obtained by Eq. (16) remains virtually unchanged by resolution changes, the other methods are fairly sensitive to them.

References

Balluch, M. G. and Haynes, P. H.: Quantification of lower stratospheric mixing processes using aircraft data, J. Geophys. Res., 102, 23487–23504, 1997.

Chipperfield, M. P.: New version of the TOMCAT/SLIMCAT offline chemical transport model: intercomparison of stratospheric tracer experiments, Q. J. Roy. Meteor. Soc., 132, 1179–1203, 2006.

Collins, W. J., Stevenson, D. S., Johnson, C. E., and Derwent, R. G.: Tropospheric ozone in a global-scale three-dimensional Lagrangian model and its response to NO$_x$ emission controls, J. Atmos. Chem., 26, 223–274, 1997.

Fairlie, T. D., Pierce, R. B., Al-Saadi, J. A., Grose, W. L., Russell III, J. M., Proffitt, M. H., and Webster, C. R.: The contribution of mixing in Lagrangian photochemical predictions of polar ozone loss over the Arctic in summer 1997, J. Geophys. Res., 104, 26597–26609, 1999.

Fukao, S., Yamanaka, M. D., Ao, N., Hocking, W. K., Sato, T., Yamamoto, M., Nakamura, T., Tsuda, T., and Kato, S.: Seasonal variability of vertical eddy diffusivity in the middle atmosphere. 1. Three-year observations by the middle and upper atmosphere radar, J. Geophys. Res., 99, 18973–18987, 1994.

Grooß, J.-U. and Russell III, J. M.: Technical note: A stratospheric climatology for O$_3$, H$_2$O, CH$_4$, NO$_x$, HCl and HF derived from HALOE measurements, Atmos. Chem. Phys., 5, 2797–2807, 2005, http://www.atmos-chem-phys.net/5/2797/2005/.

Haynes, P. and Anglade, J.: The vertical scale cascade in atmospheric tracers due to large-scale differential advection, J. Atmos. Sci., 54, 1121–1136, 1997.

Haynes, P. and Shuckburgh, E.: Effective diffusivity as a diagnostic of atmospheric transport, J. Geophys. Res., 105, 22777–22794, 2000.

Haynes, P. and Vanneste, J.: Stratospheric tracer spectra, J. Atmos. Sci., 61, 161–178, 2004.

Konopka, P., Grooß, J.-U., Günther, G., McKenna, D. S., and Müller, R.: Weak impact of mixing on chlorine deactivation during SOLVE/THESEO 2000: Lagrangian modeling (CLaMS) versus ER-2 in situ observations, J. Geophys. Res., 108, 8324, doi:10.1029/2001JD000876, 2003.

Konopka, P., Steinhorst, H.-M., Grooß, J.-U., Günther, G., Müller, R., Elkins, J. W., Jost, H.-J., Richard, E., Schmidt, U., Toon, G., and McKenna, D. S.: Mixing and ozone loss in the 1999–2000 Arctic vortex: Simulations with the three-dimensional Chemical Lagrangian Model of the Stratosphere (CLaMS), J. Geophys. Res., 109, D02315, doi:10.1029/2003JD003792, 2004.

Konopka, P., Spang, R., Günther, G., Müller, R., McKenna, D. S., Offermann, D., and Riese, M.: How homogeneous and isotropic is stratospheric mixing? Comparison of CRISTA-1 observations with transport studies based on the Chemical Lagrangian Model of the Stratosphere (CLaMS), Q. J. Roy. Meteor. Soc., 131, 565–579, 2005.

Konopka, P., Günther, G., Müller, R., dos Santos, F. H. S., Schiller, C., Ravegnani, F., Ulanovsky, A., Schlager, H., Volk, C. M., Viciani, S., Pan, L. L., McKenna, D.-S., and Riese, M.: Contribution of mixing to upward transport across the tropical tropopause layer (TTL), Atmos. Chem. Phys., 7, 3285–3308, 2007, http://www.atmos-chem-phys.net/7/3285/2007/.

Legras, B., Pisso, I., Berthet, G., and Lefèvre, F.: Variability of the Lagrangian turbulent diffusion in the lower stratosphere, Atmos. Chem. Phys., 5, 1605–1622, 2005, http://www.atmos-chem-phys.net/5/1605/2005/.

Loewenstein, M., Jost, H., Grose, J., Eilers, J., Lynch, D., Jensen, S., and Marmie, J.: Argus: a new instrument for the measurement of the stratospheric dynamical tracers, N_2O and CH_4, Spectrochim. Acta A, 58, 2329–2345, 2002.

Massie, S. T. and Hunten, D. M.: Stratospheric eddy diffusion coefficients from tracer data, J. Geophys. Res., 86, 9859–9868, 1981.

McKenna, D. S., Konopka, P., Grooß, J.-U., Günther, G., Müller, R., Spang, R., Offermann, D., and Orsolini, Y.: A new Chemical Lagrangian Model of the Stratosphere (CLaMS) 1. Formulation of advection and mixing, J. Geophys. Res., 107, 4309, doi:10.1029/2000JD000114, 2002.

Newman, P., Harris, N. R. P., Adriani, A., Amanatidis, G. T., Anderson, J. G., Braathen, G. O., Brune, W. H., Carslaw, K. S., Craig, M. S., DeCola, P. L., Guirlet, M., Hipskind, R. S., Kurylo, M. J., Küllmann, H., Larsen, N., Mégie, G. J., Pommereau, J.-P., Poole, L. R., Schoeberl, M. R., Stroh, F., Toon, O. B., Trepte, C. R., and Roozendael, M. V.: An overview of the SOLVE-THESEO 2000 campaign, J. Geophys. Res., 107, 8259, doi:10.1029/2001JD001303, 2002.

O'Rourke, J.: Computational Geometry in C, Cambridge University Press, 1998.

Prather, M. J.: Numerical advection by conservation of second-order moments, J. Geophys. Res., 91, 6671–6681, 1986.

Ray, E. A., Moore, F. L., Elkins, J. W., Hurst, D. F., Romashkin, P. A., Dutton, G. S., and Fahey, D. W.: Descent and mixing in the 1999–2000 northern polar vortex inferred from in situ tracer measurements, J. Geophys. Res., 107, 8285, doi:10.1029/2001JD000961, 2002.

Reithmeier, C. and Sausen, R.: ATTILA: Atmospheric tracer transport in a Lagrangian model, Tellus B, 54, 278–299, 2002.

Romashkin, P. A., Hurst, D. F., Elkins, J. W., Dutton, G. S., Fahey, D. W., Dunn, R. E., Moore, F. L., Myers, R. C., and Hall, B. D.: In situ measurements of long-lived trace gases in the lower stratosphere by gas chromatography, J. Atmos. Ocean. Tech., 18, 1195–1204, 2001.

Shimazaki, T. and Wuebbles, D. J.: On the theoretical model for vertical ozone density distributions in the mesosphere and upper stratosphere, Pure Appl. Geophys., 106–108, 1446–1463, 1973.

Simmons, A. J., Uppala, S. M., Dee, D., and Kobayashi, S.: ERA-Interim: New ECMWF reanalysis products from 1989 onwards, ECMWF News Lett., 110, 25–35, 2006.

Simmons, A. J., Uppala, S. M., and Dee, D.: Update on ERA-Interim, ECMWF News Lett., 111, 5 pp., 2007.

Stenke, A., Dameris, M., Grewe, V., and Garny, H.: Implications of Lagrangian transport for simulations with a coupled chemistry-climate model, Atmos. Chem. Phys., 9, 5489–5504, 2009, http://www.atmos-chem-phys.net/9/5489/2009/.

Waugh, D. W., Plumb, R. A., Elkins, J. W., Fahey, D. W., Boering, K. A., Dutton, G. S., Volk, C. M., Keim, E., Gao, R.-S., Daube, B. C., Wofsy, S. C., Loewenstein, M., Podolske, J. R., Chan, K. R., Proffitt, M. H., Kelly, K. K., Newman, P. A., and Lait, L. R.: Mixing of polar vortex air into middle latitudes as revealed by tracer-tracer scatterplots, J. Geophys. Res., 102, 13119–13134, 1997.

Webster, C. R., May, R. D., Trimble, C. A., Chave, R. G., and Kendall, J.: Aircraft (ER-2) laser infrared absorption spectrometer (ALIAS) for in-situ stratospheric measurements of HCI, N_2O, CH_4, NO_2, and HNO_3, Appl. Optics, 33, 454–472, 1994.

Wilson, R.: Turbulent diffusivity in the free atmosphere inferred from MST radar measurements: a review, Ann. Geophys., 22, 3869–3887, 2004, http://www.ann-geophys.net/22/3869/2004/.

Wohltmann, I. and Rex, M.: Improvement of vertical and residual velocities in pressure or hybrid sigma-pressure coordinates in analysis data in the stratosphere, Atmos. Chem. Phys., 8, 265–272, 2008, http://www.atmos-chem-phys.net/8/265/2008/.