TRANSPORT AND MIXING IN THE ATMOSPHERE

Peter H. Haynes
Department of Applied Mathematics and Theoretical Physics,
Wilberforce Road, Cambridge, CB3 0WA,
United Kingdom
p.h.haynes@damtp.cam.ac.uk

Abstract Transport and mixing processes in the atmosphere operate on scales from millimeters to thousands of kilometers. In certain parts of the atmosphere the large-scale quasi-horizontal flow appears to play the dominant role in transport and in the stirring process that leads ultimately to true (molecular) mixing at very small scales. Previous work in other fluid dynamical contexts such as 'chaotic advection' or 'Batchelor-regime turbulence' is therefore potentially relevant. This article reviews how, with appropriate modification, fluid dynamical insights and methods can be used, in conjunction with observational data on large-scale velocity fields or on chemical species, to quantify different aspects of transport and mixing in the atmosphere.

1. Introduction

The transport and mixing properties of the atmospheric flow are of great significance, since they play a major role in determining the distribution of atmospheric chemical species, with implications for air quality, for absorption of ultra-violet radiation (in the case of ozone) and for climate (in the case of the many chemical species that play a role in the radiation balance of the atmosphere). Indeed it is becoming common to regard the distribution of chemical species as an integral part of the climate system. Whilst air quality considerations have typically taken account of flows on the local or regional scale, it is now realised that pollutants emitted in industrialised regions can be transported thousands of kilometers and therefore that efforts to meet local quality standards often need to take account of transport by the flow on much larger scales (e.g. Akimoto 2003).

The atmosphere is conventionally divided into layers according the vertical temperature gradient. In the *troposphere* (the lowest 10 km or so of the atmosphere) the temperature decreases with height and whilst the associated density stratification is stable, the stability is relatively weak. In the *stratosphere* (roughly 10–50 km) the temperature is constant in height or increases with height and the stability is much stronger than in the troposphere. The transition from troposphere to stratosphere is called the *tropopause*. In both troposphere and stratosphere the stable density stratification, together with rotation, inhibit three-dimensionality of the flow. Thus three-dimensional turbulence tends to be confined to relatively localised regions, in the troposphere to the atmospheric boundary layer (the lowest kilometre of so of the atmosphere that is in direct contact with the Earth's surface) and to convective clouds and elsewhere in the troposphere and stratosphere to localised regions of turbulence that result from dynamical instabilities, perhaps associated with the breaking of inertia-gravity waves. There has been a tendency to regard the troposphere as relatively well-mixed in the vertical by convection, but observations of long-lived thin layers of anomalous chemical species (e.g. Newell et al. 1999) now show that, particularly in the extratropics, the time scales on which air masses experience convective events may be long – perhaps a few weeks.

Outside of regions of strong three-dimensional turbulence and on sufficiently large scales, e.g. larger than a few tens of kilometers, the atmospheric flow is quasi-horizontal, with air parcel trajectories along weakly sloping surfaces, so that horizontal displacements are generally much larger than vertical displacements. This flow has a dual character, with some aspects of its behaviour appearing organised and wave-like and other aspects exhibiting considerable nonlinearity and randomness. In the latter respect such flow might therefore be regarded as a kind of turbulence, analogous to the two-dimensional turbulence studied in idealised numerical simulations and laboratory experiments.

In considering the transport and mixing properties of flows it is useful to distinguish between three distinct processes. The first is *transport* – the bulk movement of chemical species away from locations of sources and towards regions of sinks. The second is *stirring* – deformation of geometric structure of chemical concentration fields so as to bring fluid the parcels with different chemical characteristics into closer and closer proximity. The third is *mixing* – the homogenising action of molecular diffusion on spatially varying chemical concentration fields. The importance of stirring is, of course, that by reducing length scales of chemical concentration fields it also reduces the time scale that molecular diffusion takes to act. One might therefore say that stirring governs the mixing ability of a flow and in some contexts the term 'mixing' is used to mean 'stirring'.

Three-dimensional turbulence on the one hand and the large-scale quasi-horizontal flow on the other have very different stirring and mixing properties. In three-dimensional turbulence the strong vortex stretching, and accompanying cascade of energy to small scales imply that the velocity gradient increases as horizontal length scale shrinks. It follows that deformation at a given scale is dominated by the flow at that scale. This has two important implications. The first is that that the time taken to stir a chemical field from some finite scale to the the scale at which molecular diffusion is important is independent of diffusivity (when the latter is small). The second is that the spatial structure of the concentration field is highly complex, since it mirrors the complex spatial structure of the underlying three-dimensional flow. There is a broader family of flows with these properties – we might call them '*Type I flows*' – which includes, for example, flows that are 'non-smooth' in the terminology of recent theoretical work on the effect of random flows on chemical concentration fields (e.g. Falkovich et al. 2001)

In the large-scale quasi-horizontal flow, on the other hand, the velocity field has a finite spatial scale and there is no strong increase of velocity gradients as horizontal length scale shrinks. It follows that the deformation at given scale is dominated by the flow at the large scale. This implies that the time taken to stir a chemical field from some finite scale to the the scale at which molecular diffusion is important increases logarithmically with the inverse of diffusivity (when the latter is small). It also implies that at small scales the structure of the concentration field is locally one-dimensional, consisting of filaments (in two dimensions) or sheets (in three dimensions). Again there is a broader family of flows with these properties – we might call them '*Type II flows*' – which includes, for example, 'smooth' flows or 'Batchelor turbulence' (e.g. Falkovich et al. 2001), or what are often known as 'chaotic advection' flows (e.g. Wiggins and Ottino 2004 and references therein).

The emphasis in the remainder of this article will be on transport and mixing by the large-scale quasi-horizontal atmospheric flow. It should be noted that the assertion that this is a 'Type II' flow is, to some extent, a working assumption, but is supported by basic fluid-dynamical arguments applied to strongly stratified, rapidly rotating flows (e.g. Bartello 2000, Shepherd et al. 2000).

2. Use of Atmospheric Data in Transport and Mixing Studies

Studies of transport and mixing in the atmosphere can potentially call on two types of data. The first is data on the distribution of chemical species (i.e. the concentration fields that arise as the effect of transport and mixing). Some of this data has been collected in-situ using instruments on aircraft or balloons. This has the advantage of high spatial resolution (perhaps a few tens of metres for balloon ascents, a few kilometres for horizontal aircraft sections), but, of course, is taken only along a one-dimensional section. Other such data is collected by satellite remote sensing which has lower spatial resolution (a few kilometers in the vertical, a few hundreds of kilometres in the horizontal), but provides a three-dimensional field. The second type of data are the global meteorological datasets of atmospheric dynamical quantities, including velocity fields, now routinely produced as part of the weather forecasting process. These datasets are typically at a horizontal resolution of 100 km or so. The datasets include a subtle blend of information direct from observations and information from the numerical models used to

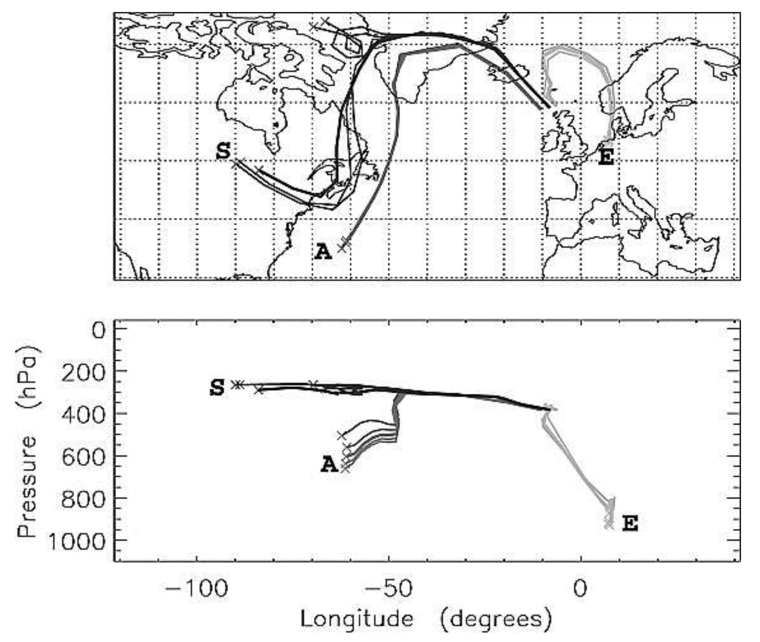

Figure 1. (From Methven et al. 2003) Three day-back trajectories from region of observation, shown in longitude-latitude in upper panel and longitude-pressure in lower panel. Copyright 2003 American Geophysical Union. Reproduced by permission of American Geophysical Union.

generate weather forecasts. The velocity datasets are now widely used as input to chemical transport models. These might be Lagrangian models, which advect large numbers of particles (or sometimes finite parcels containing reacting chemical species), or Eulerian models, which represent global fields of different chemical species. Such models have been used with great success for many different scientific purposes, e.g to interpret high-resolution one-dimensional sections of chemical species measured in observational campaigns or to predict chemical ozone destruction in the stratosphere (e.g. Waugh et al. 1994, Methven et al. 2003). Of course, this whole approach depends on the flow being Type II rather than Type I, otherwise the fact that information from spatial scales below the resolution of the datasets is not included would be a serious limitation.

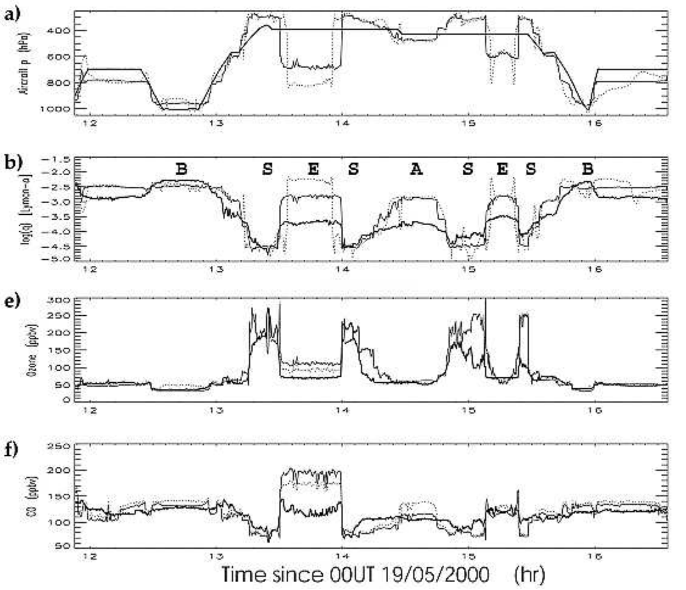

Figure 2. (From Methven et al. 2003) Time series of observations (bold lines) along the ACTO flight on 19 May 2000 compared to results from trajectory simulations. (a) Pressure and (b) log(q), where q is concentration of water in g kg^{-1}. Dotted lines show values interpolated from global meterological data to the trajectory origins at 12UT, 17 May 2000. Solid lines show an 'air-mass average' of these modeled values. (c) and (d) not shown. (e) Ozone and (f) carbon monooxide concentrations. Dotted line is air-mass average at origin of trajectories. Solid is prediction of chemical model integrated along the trajectory. Copyright 2003 American Geophysical Union. Reproduced by permission of American Geophysical Union.

The Methven et al. (2003) paper gives a nice example of the sort of work that it is now possible using a combination of chemical data (in this case from aircraft) and Lagrangian calculations using the global velocity fields. The focus of the paper is a campaign (ACTO) in which a research aircraft made chemical and meteorological measurements to the northwest of Scotland. Figure 1 shows back-trajectory calculations that suggest that air sampled in the upper troposphere in this region on one particular day of the campaign converged from three distinct locations. Some of the back trajectories originate in region A, the mid-troposphere in Eastern Atlantic, some in region S, the stratosphere over central Canada, and some in region E, the lower troposphere over central Europe. Air from each of these regions has a different chemical signature. That from region A is moist and relatively low in ozone (O_3) and carbon monooxide (CO). That from region S is dry, high in O_3 and low in CO. That from region E is similar to that from A in that it is moist (both regions are in the troposphere) but different in that it is relatively polluted and therefore high in CO and in O_3 (some of which is likely to have formed through photochemical production as air moves from region E to the region of measurement). Comparison with the detailed chemical fields measured by the aircraft, as shown in Fig. 2, shows that the positions and chemical characteristics of the different air masses are generally well predicted by the back trajectory calculation, even, in many cases, down to small-scale features. See Methven et al. (2003) for further details and discussion.

3. Transport and Stirring

Given that the large-scale atmospheric flow is quasi-horizontal, models of chaotic advection in two-dimensional incompressible flow are potentially relevant. Such models and analogous dynamical systems show that transport and stirring is often highly inhomogeneous, with regions of strong stirring (i.e. strong stretching) separated by barrier regions in which stirring is weak and across which there is no, or relatively little, transport (e.g. Meiss 1992, Wiggins 1992). Many detailed investigations over the last 20 years or so have shown how the transport and stirring properties of simple flows change with the parameters defining that flow. The typical picture is as follows. For a steady incompressible two-dimensional flow the dynamical system describing particle advection is integrable. Particles move along streamlines, i.e. contours of the streamfunction, and these streamlines, which are invariant curves of the dynamical system, act as perfect barriers to transport. If a small time-periodic perturbation is added to the flow then, consistent with the

KAM theorem, many (in fact 'almost all' in the limit of a vanishingly small perturbation) invariant curves survive. The same is probably true for relatively large classes of quasi-periodic perturbations. Between the surviving invariant curves there are thin regions in which particle trajectories are chaotic (and stirring is therefore strong). These chaotic regions increase in thickness as the amplitude of the perturbation increases, with the thickest regions usually centred on the location of streamlines of the unperturbed flow that pass through hyperbolic stagnation points. For small perturbations the geometry of surviving invariant curves is similar to that of the streamlines of the steady flow and many features of the pattern of transport in the steady flow therefore persist to the perturbed flow. For example, two disjoint regions may be separated by many streamlines of the steady flow and by many invariant curves of the weakly perturbed flow, so that transport between those two regions is forbidden. However as the size of the perturbation increases more and more invariant curves break, so that ultimately the last invariant curve separating the two disjoint regions breaks and transport between those regions is allowed.

This sort of behaviour is illustrated in Fig. 3, which shows Poincaré sections for simple time-periodic flows, with the amplitude of the time-dependent perturbation measured by a parameter ϵ_2. For more details see Shuckburgh and Haynes (2003). For $\epsilon_2 = 0.0125$ and $\epsilon_2 = 0.025$ there is a perfect central barrier to transport roughtly indicated by the grey curve. For $\epsilon_2 = 0.05$ and $\epsilon_2 = 0.075$ this central barrier has disappeared. Nonetheless, for $\epsilon_2 = 0.05$ some memory of the barrier persists and transport between the two halves of the flow domain is relatively slow, as can be seen by (at finite time) the modest intersection of the two Poincaré sections. Again this persistence of a partial barrier after the last invariant curve has disappeared is a generic feature of such systems.

In the real atmosphere observations and models show similarly inhomogeneous transport and stirring. In the winter stratosphere, for example, the edge of the winter polar vortex seems to form a transport barrier which is not perfect, but across which transport is very weak. Outside the vortex, in mid-latitudes, and to a lesser extent inside the vortex, there are regions of strong stirring (e.g. Joseph and Legras 2002 and references therein). Lower in the atmosphere the subtropical jet seems to form a similar, if less effective, barrier to quasi-horizontal transport which separates the upper troposphere and the lowest part of the stratosphere. Once again, poleward and equatorward of the subtropical jet there are regions of strong stirring (e.g. Haynes and Shuckburgh 2000b). This sort of spatial structure, with transport barriers and stirring regions, is found in a wide class of geophysical flows.

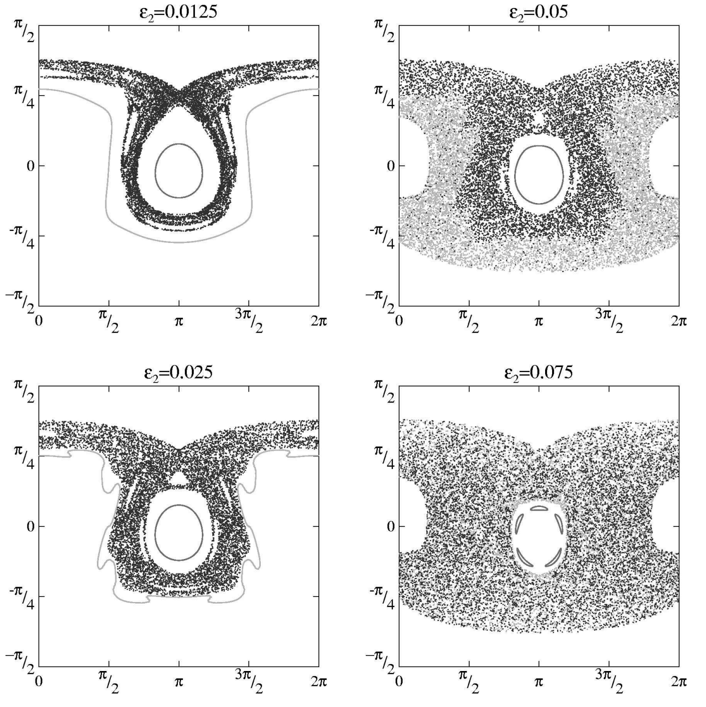

Figure 3. (From Shuckburgh and Haynes 2003) Poincaré sections (10000 periods) for a simple time periodic flow, with ϵ_2 measuring the amplitude of the time-dependent part of the flow. Three different Poincaré sections are plotted in each case. The limited intersection between two such sections for $\epsilon_2 = 0.05$ shows that whilst a 'last' invariant curve has broken, significant inhibition of transport remains. Copyright 2003 American Institute of Physics. Reproduced by permission of American Institute of Physics.

On a purely diagnostic level, quantifying the transport and stirring properties of an atmospheric flow, as defined by global meteorological datasets or by numerical models, poses the same challenges as quantifying transport and stirring (or 'mixing') in any chaotic dynamical system. Such quantification is of practical important, e.g to compare effectively the transport and stirring structure of velocity fields in different observational datasets or generated by different numerical models. Various mathematical tools originating in dynamical systems theory have been studied in this context. One example is the method of 'lobe dynamics' which calculates fluxes across control surfaces built up of parts of stable and unstable manifolds of particular hyperbolic trajectories (which in the particular case of time-periodic flow are usually hyperbolic fixed points of the corresponding Poincaré map) (Wiggins 1992, Malhotra and Wiggins 1998). Koh and Plumb (2000) have applied this method to quantify transport across the polar vortex edge in a simple model and find that it gives a misleading estimate of such transport. The reason seems to be that the type of control surface that is usually constructed in lobe dynamics does not usefully correspond to the vortex edge (see Joseph and Legras 2002 for further discussion). The same difficulty is likely to arise in any attempt to use lobe dynamics to quantify transport across a disturbed jet (and many atmospheric transport problems are of this nature). It may be that these difficulties may be circumvented by a more ingenious choice of control surface. It should also be noted that lobe dynamics is more successful and, potentially, practically useful for other geophysical flows, e.g. the two-gyre oceanic flow studied by Coulliette and Wiggins (2001).

A completely different method to quantifying transport that has been applied to realistic atmospheric flows is to calculate 'effective diffusivity' – a measure of the geometry of a test tracer that is advected by the flow (Haynes and Shuckburgh 2000a,b). This method is apparently successful in identifying transport barriers associated with jets or edges of vortices and in calculating their strength. Shuckburgh and Haynes (2003) examine the relation of effective diffusivity to other measures of transport and stirring.

A more fundamental question concerns the apparent close similarity between the transport and stirring structure (barriers and stirring regions) seen in time-periodic flows and those observed in geophysical flows, including realistic atmospheric flows. The fact is that in the former the flow is imposed in advance (the model is 'kinematic'), whereas in the latter the flow is the solution of a set of dynamical equations. In large-scale atmospheric and oceanic flows, those dynamical equations can be expressed to good approximation as the conservation of potential vorticity (PV) following the fluid motion (with non-conservation associated with frictional or diabatic effects), together with an invertibility relation that determines all other flow quantities instantaneously from the potential vorticity. Thus the transport and stirring properties of a flow are, because of transport and stirring of potential vorticity, inextricably linked to the dynamics. Study of the transport and stirring structure of given flows therefore poses a 'dynamical consistency' problem – it is not self-consistent to ignore the effect of the transport and stirring on the dynamics. Brown and Samelson (1994) argue in the case of two-dimensional PV conserving flow that one implication of dynamical consistency is that chaotic particle trajectories are forbidden, though it may be possible to escape this constraint if PV is not perfectly conserved.

One argument for the relevance of the kinematic models is that the potential vorticity field arising from the transport and stirring is in many cases dominated by small spatial scales and the smoothing nature of the inversion operator is such that the dynamical effect is weak. However it is difficult to see that the dynamical effect will always be weak and it seems likely that in some aspects of transport and mixing behaviour, e.g. the persistence versus destruction of transport barriers as perturbation amplitude is increased, the dynamical effect will be significant.

It seems extremely difficult to formulate any analytic or semi-analytic model that allows an interesting combination of dynamical consistency and chaotic transport. All attempted formulations have, to the author's knowledge, ended up considering cases where the transport (of potential vorticity) is for one reason or another uncoupled from the dynamics. An alternative approach is to consider suitable numerical experiments. Poet (2004) reports a comparison between a kinematic model of a meandering jet in a channel previously considered by Pierrehumbert (1991) and others and a corresponding dynamically consistent model based on the two-dimensional vorticity equation on a β-plane. The 'correspondence' is that if the dynamical equations are linearised about a background unidirectional flow along the channel the flow in the dynamically consistent model would be precisely that in the kinematic model.

The kinematic model is different to that for which results are shown in Fig. 3, but has the same generic behaviour, where for small time-dependence many invariant curves persist and transport from one half of the channel to the other is forbidden. Then at a critical value of the parameter ϵ governing the amplitude of the time-dependent component of the flow, the 'last' invariant curve breaks and transport from one half of the channel to the other is permitted. This is indicated in Fig. 4 (curve A) by the average transit time (from one chosen region to another) which is essentially infinite for $\epsilon < 0.34$. The slow decrease of the transit time as ϵ increases beyond 0.34 indicates that significant inhibition of transport remains (analogous to the behaviour seen for $\epsilon_2 = 0.05$ in Fig. 3). The results for the dynamically consistent model are shown as curve B. Here it is difficult to set up conditions where any transport barrier is perfect. The large average transit time for $\epsilon < 0.44$ may therefore be taken to indicate a strong barrier. In this sense the strong barrier persists to a larger value of ϵ in the dynamically consistent model than in the kinematic model – dynamical consistency apparent strengthens the barrier effect in that sense. A further important difference is that for $\epsilon < 0.44$, once the barrier has broken, the transit time is immediately significantly reduced. There is therefore no persistence of a partial barrier as seen in the kinematic case. The reason is that, in the dynam-

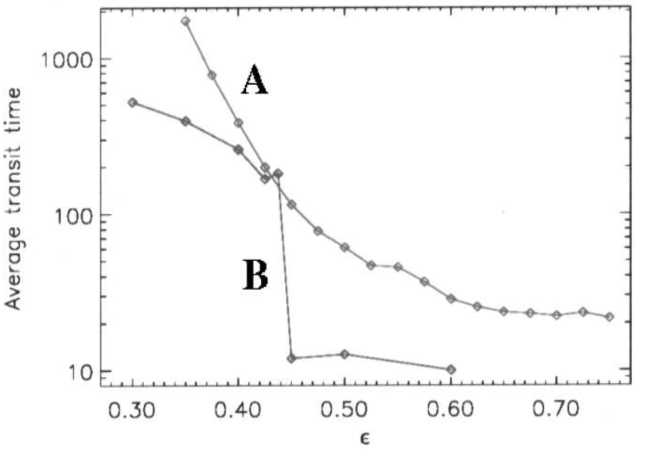

Figure 4. Transit time across central barrier as function of ϵ, amplitude of time-dependent component of flow. The upper curve A is for the kinematic model and the lower curve B is for the dynamically consistent model. See text for further details. Reproduced from Poet(2004). Copyright D.A. Poet 2004.

ically consistent case, once the central barrier is broken the potential vorticity is stirred throughout the flow domain and the character of the flow changes drastically as a result.

4. Stirring and Mixing

Mixing in the atmosphere could be achieved in laminar flow, but in practice is likely to be enhanced through intermittent encounters of air parcels with three-dimensional turbulence. This enhancement might be substantial in the troposphere, where convection is relatively common, is likely to be much weaker in the lower stratosphere and then to increase again in the upper stratosphere and mesosphere, where three-dimensional turbulence associated with the breaking of gravity waves is more widespread. One outstanding question is how to estimate the overall mixing effects of these encounters. One approach is by direct observation of the turbulent events themselves (e.g. Alisse et al. 2000), but the spatial intermittency and strong geographical and seasonal variation make it difficult to integrate such observations to give a quantitative estimate of overall mixing effect.

A different approach is to argue that the chemical concentration fields in the atmosphere arise from a combination of advection and mixing effects. If the former is dominated by the large-scale flow, then it can to some extent be estimated using the velocity fields from global meteorological datasets. This gives the possibility of an 'inverse' calculation, taking the observed distributions of chemical species and trying to deduce what quantitative representation of mixing is consistent with those observations and with the (known) large-scale flow.

Predictions of the three-dimensional structure of chemical fields in necessary detail, given the large-scale flow, are beyond current computational resources (whatever mixing processes are assumed). However, there are theories for stirring and mixing of chemical species in large-scale 'Type-II' flows that offer an alternative approach to full three-dimensional computation.

A long-standing theoretical approach has been to consider the interaction between stretching and diffusion in a flow that is a random function of time and a linear function of space. This theory has undergone significant developments recently (e.g. Antonsen et al. 1996, Falkovich et al. 2001). Whilst the theory has serious limitations in the initial-value problem (i.e. free decay from a specified initial field for the chemical concentration) (Sukhatme and Pierrehumbert 2002, Fereday and Haynes 2004) these limitations do not seem to extend to the forced problem (where the chemical concentration is maintained by some source/sink distribution) and potentially it offers a quantitative theory for the chemical concentration field that arises from a specified combination of large-scale flow and small-scale mixing processes. The standard theories need some extension to be applied to the atmosphere, but this extension is relatively straightforward. On this basis Haynes and Vanneste (2004) have examined the effect of diffusivity (as a representation of small-scale turbulent mixing processes) on the spatial structure, specifically the wavenumber spectrum, of long-lived chemical species in the lower stratosphere. The simplest theoretical prediction is that the spectrum has a k^{-1} range, where k is wavenumber at small k (but larger than the forcing wavenumber) and a steeper spectrum at larger k where mixing effects become important. The results suggest that, for plausible stratospheric forcing scales, a clear k^{-1} regime would be seen only if the diffusivity D were similar in value to molecular diffusivity at about $10^{-4}\,\mathrm{m^2s^{-1}}$, implying little turbulent enhancement of mixing. If D were $10^{-2}\,\mathrm{m^2s^{-1}}$, on the other hand, which corresponds to the value estimated by the Alisse et al. (2000) study, then on scales of a few kilometers the form of the spectrum is rather close to k^{-2}, which happens to coincide with observational estimates by Sparling and Bacmeister (2001). This is certainly not the last word, however. More attention needs to be paid to the effects of spatial and temporal organisation of the stratospheric flow (the persistent stratospheric polar vortex, for example), which is not properly incorporated by the theories. Furthermore, calculations by Legras et al. (2003) using a different approach suggest that a larger value of D gives a better match to observations. Finally, similar studies are needed of the tropopause region where mixing processes are ill-quantified and may have significant effects on chemical distributions.

References

H. Akimoto, Global air quality and pollution, *Science*, 302, pp.1716-1719, 2003.

J.-R. Alisse, P.H. Haynes, J. Vanneste, C. Sidi, Quantification of stratospheric mixing from turbulence microstructure measurements, *Geophys. Res. Lett.*, 27, pp.2621-2624, 2000.

T.M. Antonsen, Z. Fan, E. Ott, E. Garcia-Lopez, The role of chaotic orbits in the determination of power spectra of passive scalars, *Phys. Fluids*, 8, pp.3094–3104, 1996.

P. Bartello, Using low-resolution winds to deduce fine structure in tracers, *Atmos. -Ocean*, 38, pp.303–320, 2000.

M.G. Brown, R.M. Samelson, Particle motion in incompressible, inviscid vorticity-conserving fluids, *Phys. Fluids*, 6, pp.2875–2876, 1994.

C. Coulliette, S. Wiggins, Intergyre transport in a wind-driven, quasigeostrophic double gyre: An application of lobe dynamics, *Nonlinear Proc. Geophys.*, 7, pp.59–85, 2001.

G. Falkovich, K. Gawedzki, M. Vergassola, Particles and fields in fluid turbulence, *Revs. Mod. Phys.* 73, pp.913–975, 2001.

D.R. Fereday, P.H. Haynes, Scalar decay in two-dimensional chaotic advection and Batchelor-regime turbulence, *Phys. Fluids*, 16, pp.4359–4370, 2004.

P.H. Haynes, E.F. Shuckburgh, Effective diffusivity as a diagnostic of atmospheric transport. Part I: stratosphere, *J. Geophys. Res.* 105, pp.22777-22794, 2000a.

P.H. Haynes, E.F. Shuckburgh, Effective diffusivity as a diagnostic of atmospheric transport. Part II: troposphere and lower stratosphere, *J. Geophys. Res.*, 105, pp.22795-22810, 2000.

P.H. Haynes, J. Vanneste, Stratospheric tracer spectra, *J. Atmos. Sci.*, 61, pp.161–178, 2004.

B. Joseph, B. Legras, Relation between kinematic boundaries, stirring and barriers for the Antarctic polar vortex, *J. Atmos. Sci.* 59, pp.1198–1212, 2002.

T.-Y. Koh, R.A. Plumb, Lobe dynamics applied to barotropic Rossby-wave breaking, *Phys. Fluids* 12, pp.1518-1528, 2000.

B. Legras, B. Joseph, F. Lefévre, Vertical diffusivity in the lower stratosphere from Lagrangian back-trajectory reconstructions of ozone profiles, *J. Geophys. Res.*, 108:D4562 – doi:10.1029/2002JD003045, 2003.

N. Malhotra, S. Wiggins, Geometric structures, lobe dynamics, and Lagrangian transport in flows with aperiodic time-dependence, with applications to Rossby wave flow, *J. Nonlinear Sci.*, 8:401–456, 1998.

J.D. Meiss, Symplectic maps, variational principles and transport, *Rev. Mod. Phys.*, 64, 795–848, 1992.

J. Methven, S.R. Arnold, F.M. O'Connor, H. Barjat, K. Dewey, J. Kent, N. Brough, Estimating photochemically produced ozone throughout a domain using flight data and a Lagrangian model, *J. Geophys. Res.*, 108:D4271–doi:10.1029/2002JD00 2955, 2003.

R.E. Newell, V. Thouret, J.Y.N. Cho, P. Stoller, A. Marenco, H.G. Smit, Ubiquity of quasi-horizontal layers in the troposphere, *Nature*, 398, pp.316–319, 1999.

R.T. Pierrehumbert, Chaotic mixing of tracer and vorticity by modulated traveling Rossby waves, *Geophys. Astrophys. Fluid Dynamics*, 58, pp.285–319, 1991.

D.A. Poet, *Dynamical consistency, barriers to transport and mixing, in chaotic advection flows*, Ph.D. thesis, University of Cambridge, pp.207, 2004.

T.G. Shepherd, J.N. Koshyk, K. Ngan, On the nature of large-scale mixing in the stratosphere and mesosphere, *J. Geophys. Res.*, 105:, pp.12433–12446, 2004.

E.F. Shuckburgh, P.H. Haynes, Diagnosing transport and mixing using a tracer-based coordinate system, *Phys. Fluids*, 15, pp.3342-3357, 2003.

L.C. Sparling, J.T. Bacmeister, Scale dependence of tracer microstructure: PDFs, intermittency and the dissipation scale, *Geophys. Res. Lett.*, 28: pp.2823–2826, 2001.

J. Sukhatme, R.T. Pierrehumbert, Decay of passive scalars under the action of single scale smooth velocity fields in bounded two-dimensional domains: From non-self-similar probability distribution functions to self-similar eigenmodes, *Phys. Rev. E.*, 66:056302-1–056302-11, 2002.

D.W. Waugh, et al. Transport of material out of the stratospheric Arctic vortex by Rossby wave breaking, *J. Geophys. Res.*, 99, pp.1071–1088, 1994.

S. Wiggins, *Chaotic transport in dynamical systems*, New York, Spring-Verlag, pp.301, 1992.

S. Wiggins, J.M. Ottino, Foundations of chaotic mixing, *Phil. Trans. Roy. Soc. Lond. A*, 362, pp.937–970, 2004.