

Technical note: a new method for the Lagrangian tracking of pollution plumes from source to receptor using gridded model output

R. C. Owen^{1,2} and R. E. Honrath^{1,2,3}

¹Department of Geological and Mining Engineering and Sciences, Michigan Technological University, USA

²Department of Civil and Environmental Engineering, Michigan Technological University, Houghton, Michigan, USA

³Atmospheric Sciences Program, Michigan Technological University, USA

Received: 23 July 2008 – Published in Atmos. Chem. Phys. Discuss.: 3 November 2008

Revised: 9 March 2009 – Accepted: 9 March 2009 – Published: 8 April 2009

Abstract. Lagrangian particle dispersion models (LPDMs) are powerful and popular tools used for the analysis of atmospheric trace gas measurements. However, it can be difficult to determine the transport pathway of emissions from their source to a receptor using the standard gridded model output, particularly during complex meteorological scenarios. In this paper we present a method to clearly and easily identify the pathway taken by only those emissions that arrive at a receptor at a particular time, by combining the standard gridded output from forward (e.g., concentration) and backward (e.g., residence time) LPDM simulations. By comparing the pathway determined from this method with particle trajectories from both the forward and backward models, we show that this method successfully restores much of the Lagrangian information that is lost when the data are gridded. A sample analysis is presented, demonstrating that the source-to-receptor pathway determined from this method is more accurate and easier to use than existing methods using standard LPDM products (gridded fields of, e.g., concentrations and residence time). As demonstrated in an evaluation and an example application, the method requires agreement between the transport described by the forward and backward simulations and thus provides a means to assess the quality and reversibility of the simulation. Finally, we discuss the potential for combining the backward LPDM simulation with gridded data from other sources (e.g., chemical transport models) to obtain a Lagrangian sampling of the air that will eventually arrive at a receptor. Based on the advantages presented here, this new method can complement or even replace many of the standard uses of backward LPDM simulations.

1 Introduction

The transport experienced by a plume of emissions can have a significant influence on its chemical composition. Dry deposition, which is an important removal mechanism for many trace gases, occurs in the boundary layer (BL). Significant wet deposition is often associated with strong uplift from the BL into the free troposphere (FT) (e.g., Stohl et al., 2002b). For example, soluble species, such as HNO₃, can be removed during this uplift. Once in the FT, however, the chemical composition of an air mass is more dependent on photochemistry and mixing (e.g., Methven et al., 2003). Thus, knowing the amount of time an air mass spends in the BL, the timing and location of uplift, the time spent in the FT, and the relative amounts of mixing during these processes is essential to a complete understanding of the chemical transformations occurring in an air mass. As a result, determining these transport characteristics for a plume of emissions as it travels from its source to a downwind sample location has been an important part of many measurement efforts (e.g., Rex et al., 1998; Stohl and Trickl, 1999; Trickl et al., 2003).

The atmospheric transport pathway through which emissions traveled to a downwind receptor is often deduced with Lagrangian models, either trajectories or Lagrangian particle dispersion models (LPDMs). Trajectories remain popular because they are easy to use, but they are limited by their inability to describe the deformation of an air mass and the concentration gradients of chemical trace substances in the atmosphere (Stohl et al., 2002a; Methven et al., 2006). LPDMs are superior because they address both these issues (e.g., Han et al., 2005), but they also have shortcomings. Their output is more complex than that of trajectory models and much of the Lagrangian information is lost in the process



Correspondence to: R. E. Honrath
(reh@mtu.edu)

of calculating concentrations on a Eulerian-type output grid. While some work has been done to simplify the LPDM output (e.g., Stohl et al., 2002a), there remains a need for new products to succinctly describe LPDM output. Additionally, there are no methods available to retrieve the Lagrangian information that is lost when the output is gridded, short of saving particle trajectories. While particle trajectories provide useful information, they are typically not saved during LPDM simulations because they increase the time it takes to run and process a simulation and require large amounts of storage and memory. Particle trajectories also add a level of complexity in interpreting the output and are generally used only in advanced LPDM studies (e.g., Stohl et al., 2004).

Traditionally, studies that use LPDMs to perform a detailed analysis of the transport of emissions to a particular receptor use both forward simulations (simulations of atmospheric concentrations resulting from an emissions field) and backward simulations (simulations of the upwind transport of air ultimately reaching a receptor), but present them as individual products. For example, Stohl et al. (2003) used forward simulations to understand the large scale transport from North America to aircraft-based sample locations over the North Atlantic and Europe, while backward simulations were used to determine the age distribution of the trace substance in the receptor cells, to determine the specific source regions contributing to the trace substance enhancements in the receptor cells, and to determine the transport pathway to the receptor. They assumed that the backward plume matched the pathway taken by emissions, which was reasonable only because this particular transport experienced very little deformation. Owen et al. (2006) analyzed the forward and backward products together, by presenting snapshots of the two simulations side-by-side, but detailed analysis of transport was limited to only those backward simulations that experienced little deformation between the source and receptor. Despite these applications, there remains a disjunction between the information provided by forward and backward LPDM simulations.

In this paper, we present and evaluate a new method, the product of which we call a folded retroplume. The folded retroplume addresses two of the shortcomings of the LPDM by simplifying the LPDM output and allowing the retrieval of some of the Lagrangian information that is lost in the process of calculating gridded (Eulerian) output fields. The purpose of the folded retroplume is to provide a way to efficiently and accurately determine the transport pathway of emissions to a receptor, highlighting only those emissions that arrive in the receptor cell at the time of interest, using standard gridded output fields from an LPDM. As we show below, this can be accomplished by convolving the standard output from a forward model simulation with that from a backward model simulation, bringing the information from the forward and backward models together in such a way that even complex transport scenarios can be analyzed. When used in this way, the forward model describes the amount of a trace substance

in the atmosphere, while the backward model describes how much of this trace substance will arrive in the receptor cell at a given time. The folded retroplume is easier to use and more accurate than using standard gridded LPDM products alone. Additionally, the method is superior to similar methods available with trajectories, as it retains the advantages of LPDMs, e.g., the ability to describe dispersion and to survey information about the relative concentration of the trace substance along the transport pathway. Since the method requires an agreement in the transport described by the two model simulations, it also allows for an assessment of the quality and reversibility of the simulation. We also introduce similar uses of the backward model with alternate Eulerian fields that describe the state of the atmosphere (e.g., output from a chemical transport model) to determine physical and chemical properties of an air mass as it travels toward a receptor.

Below, we provide a method overview (Sect. 2), evaluation (Sect. 3), and example (Sect. 4) that are based on the LPDM FLEXPART (Stohl et al., 2005), one of the more popular LPDMs in use today. Although the presentation is somewhat specific to this model, the method should be valid for any LPDM with an appropriately employed backward mode, because all LPDMs have the property that they are essentially self-adjoint, i.e., the backward mode only requires the reversal of the direction of advection to give the transport sensitivity for a receptor cell (Seibert and Frank, 2004).

2 Method overview

We begin with a brief outline of the formulation of the model output and the folded retroplume, followed by a simple case that illustrates the folded retroplume. We rely heavily on the model theory presented by Stohl et al. (2005) and Seibert and Frank (2004) and refer the reader to these sources for more detailed reviews of LPDM theory and technical descriptions of LPDM operational details. We also recommend Flesch et al. (1995) and Lin et al. (2003) for additional information on backward LPDM modeling and Errico (1997) and for additional background on general adjoint model theory.

2.1 Formulation of the model output and folded retroplume

In this section, we review the calculation of the gridded model output, starting with the forward mode, and provide several formulations involving the folded retroplume. The calculations presented in this section are limited to instantaneous model output. (The use of averaged model output can complicate the interpretation of folded retroplumes and is thus discussed in Sect. 3.3.)

2.1.1 Standard output from the forward mode

In the forward mode, particles are released at the source and then transported forward in time, according to the mean and turbulent wind components (Stohl et al., 2005). The mass of each particle at the time of emission is based on the strength of the source. Concentrations are calculated by summing the mass of all the particles that reside in each grid cell. We first consider only a puff of emissions and focus on one downwind grid cell (j) at a single model time (t). The instantaneous gridded concentration (c) from a puff of emissions released at time t_0 is thus:

$$c_{j,t} = \frac{1}{V_j} \sum_{V_j} m_{j,t_0} \cdot p_{j,t} \cdot f_j, \quad (1)$$

where V_j is the volume of the cell and the summation is over all particles that reside in V_j at time t . m_{j,t_0} is the initial mass of the particle, f_j is the sampling kernel, which can be used to distribute the mass of the particle across multiple grid cells, and $p_{j,t}$ is the transmission function, which describes the percentage of the particle mass remaining from removal processes (see Stohl et al., 2005 and Seibert and Frank, 2004 for a more complete description of these terms). In order to calculate the mass mixing ratio, the concentration is first divided by the local air density (from the meteorological data). The volume term from the concentration cancels with the volume term from the local air density, leaving the mass of air in the cell ($m_{j,\text{air}}$) and the summation of the mass of the particles in the cell. The volume mixing ratio (χ) is obtained by multiplying this value by the ratio of the mean molecular mass of air (M_{air}), to that of the trace substance being modeled (M_{tr}), giving the volume mixing ratio:

$$\chi_{j,t} = \left(\frac{M_{\text{air}}}{M_{tr}} \right) \left(\frac{1}{m_{j,\text{air}}} \right) \sum_{V_j} m_{j,t_0} \cdot p_{j,t} \cdot f_j. \quad (2)$$

Mixing ratios are saved at each time for each grid cell in the model domain, giving a 3-dimensional matrix (χ_t) of the volume mixing ratios for all j (i.e., each x , y , and z component).

2.1.2 Standard output from the backward mode

In the backward mode, particles are initiated in a single receptor volume (j') over a short interval (t' , the arrival time) and transported backward in time by reversing the direction of the mean wind. The mass of each particle is normalized by the total mass released in the receptor (m_{tot}), giving each particle units of mixing ratio, such that each particle represents one part of the air in the receptor at the release (i.e., arrival) time and the distribution of the particles indicates the location of the receptor air at each upwind time. The backward output is gridded by summing these mixing ratios in each cell, giving the sensitivity (S) of the receptor to the mass present in the upwind cell:

$$S_{j,t,(j',t')} = \sum_{V_j} \frac{m_{j,t'}}{m_{\text{tot}}} \cdot p_{j,t'} \cdot f_j, \quad (3)$$

Again, the output is saved at each time for each grid cell, giving a 3-dimensional matrix ($\mathbf{S}_{t,(j',t')}$) of the sensitivity for all j . The output of the backward mode is referred to as the sensitivity plume, or the retroplume.

2.1.3 The folded retroplume – combining model output to determine the source-to-receptor transport pathway

The folded retroplume at time t is the Hadamard (or entry-wise) product of the mixing ratio matrix (χ_t) from the forward mode and the sensitivity matrix ($\mathbf{S}_{t,(j',t')}$) from the backward mode. In terms of an individual cell, the mixing ratio (Eq. 2) indicates the amount of emitted trace substance in the given cell, and the sensitivity (Eq. 3) indicates the amount of air in the cell that will be transported to the receptor. By multiplying the two values, we can determine the amount of the trace substance in the cell that will eventually arrive at the receptor (j') at the arrival time (t'). Note the units for this operation. We begin with the volume mixing ratio, with units of parts of trace substance per parts of air in the cell. This is multiplied by the sensitivity, with the units of parts of air in the cell per parts of air in the receptor. The resulting product calculated for a specific cell j has units of volume mixing ratio, and indicates the portion of the mixing ratio in the receptor at t' (the sensitivity plume arrival time) resulting from the transport of trace substance through cell j at time t . That is, the units are parts of trace substance in the cell per part of air at the receptor. As this mixing ratio results from only a part of the sensitivity field (the individual upwind cell considered here), we call the product the partial mixing ratio (PMR):

$$\text{PMR}_{j,t,(j',t')} = S_{j,t,(j',t')} \cdot \chi_{j,t}, \quad (4)$$

where t is the model time. The PMR will clearly be small or zero when there is either little of the trace substance in a cell or small sensitivity. Conversely, if there is a significant amount of trace substance in a cell and a large sensitivity, then the PMR will also be large, indicating the location of trace substance that travels from the source to the receptor. The 3-dimensional matrix $\mathbf{PMR}_{t,(j',t')}$ indicates the distribution at time t of the trace substance that will ultimately arrive at the receptor at time t' , while the matrix $\mathbf{PMR}_{(j',t')}$ at multiple times shows the transport pathway of the trace substance between the source and receptor.

Up to this point, we have only considered a puff of emissions in the forward model, which is not the typical model situation. Normally, emissions are continuously released into the forward simulation and each particle is carried in the model for a set number of days and then dropped. Thus, the mixing ratio from the forward model ($\chi_{j,t}$) consists of particles released over a range of times and can be divided into age classes, according to the length of time the particles have been in the model (the age of the particles). If the age of the trace substance is not taken into account when computing

the PMR then the folded retroplume calculation will include particles that would be dropped from the forward simulation before they would be transported to the receptor. To avoid this, the age of forward model trace substance in Eq. (4) ($\chi_{j,t}$) must be less than the time difference between the release time in the backward model (that is, the arrival time) and the sample time (t). If we have a forward simulation that carries particles for A_f days and a backward simulation with an arrival time of t' , then the PMR at some intermediate time (t) should be:

$$\text{PMR}_{j,t,(j',t')} = S_{j,t,(j',t')} \cdot \sum_{i=0}^{A_f-(t'-t)} \chi_{j,t,i} \quad (5)$$

where i indicates the available age classes in days from the forward model.

The PMRs at any upwind time may be summed over the model domain to determine the mixing ratio from all (appropriately aged) emissions that are present in the model at that time. If no more emissions are added to the atmosphere between that time and arrival time, then this sum would be equal to the mixing ratio in the receptor (j') at the arrival time (t'). Thus, we call this sum the upwind mixing ratio (UMR):

$$\text{UMR}_{t,(j',t')} = \sum_j \text{PMR}_{j,t,(j',t')} \quad (6)$$

The UMR is equivalent to a sensitivity-weighted average of the upwind mixing ratio field and can increase or decrease, depending on the relative rates of emission and removal processes. For example, if emissions are added to the atmosphere between time steps and no removal processes are considered, the change in the UMR from time t to $t+1$ should be

$$\text{UMR}_{t+1,(j',t')} = \text{UMR}_{t,(j',t')} + \sum_j S_{j,t,(j',t')} \cdot E_{j,t} \quad (7)$$

where $E_{j,t}$ are the emissions released into the model at time t and $\sum_j S_{j,t,(j',t')} \cdot E_{j,t}$ is the so called source contribution (Stohl et al., 2003). However, if no removal processes are considered and if no emissions are added to the atmosphere after time t (or if no emissions are added in areas with sensitivity – the region where the plume is located), then the UMR should remain constant.

The UMRs therefore provide a means to evaluate the evolution of the mixing ratio of the receptor air during transport. For instance, the timing and location of wet removal could be determined by comparing the UMRs from two folded retroplumes, one computed from forward and backward simulations with no wet removal and one computed from forward and backward simulations that include wet removal. Section 5 will discuss other potential applications using the UMRs derived from folding a backward simulation with mixing ratio fields from alternate sources.

2.2 Illustrative case

Here we present the application of our method to a simple case in order to illustrate the folded retroplume. The case is based on a puff of CO emissions released into the forward model from the Boston area, into the box bounded by 41–43° N latitude and 73–75° W longitude, from the surface up to a height of 250 m a.s.l. Emissions were released over a 1-h period, from 15:00–16:00 UTC on 14 May 2005 and were based on the EDGAR Fast Track 1999 inventory (Olivier et al., 1996). The backward simulation was initiated at the Pico Mountain observatory, located on the Azores Islands in the Central North Atlantic Ocean, into the box bounded by 38.5–39.0° N latitude 28.5–28.0° W longitude, from an altitude of 2000–2250 m a.s.l. Particles for the backward simulation were also released over a 1-h period from 00:30–01:30 UTC on 19 May 2005.

FLEXPART version 6.2 was used, driven with data from the European Centre for Medium Range Weather Forecasts (ECMWF) (ECMWF, 2005) with 1°×1° horizontal resolution, 60 vertical levels and a temporal resolution of 3 h, using meteorological analyses at 00:00, 06:00, 12:00 and 18:00 UTC and ECMWF 3-h forecasts at intermediate times (03:00, 09:00, 15:00 and 21:00 UTC). The output was saved with a grid size of 0.5°×0.5° in the horizontal and 250 m in the vertical, from 0–7000 m a.s.l. The sampling kernel was turned off and instantaneous fields were saved (see Sect. 3.3 for a description of these model settings and more details on their impact on the folded retroplume). 500 000 particles were used for the forward simulation and 2000 particles were used for the backward simulation, resulting in a total of 670 forward particles and 634 backward particles that successfully travel between the source and receptor cells.

Figure 1a and b shows the plan view and longitude-height cross section of the CO plume 1.5 days after the forward-model puff release. Figure 1c and d shows plan view and longitude-height cross section of the sensitivity plume 3 days upwind of the release at the receptor (and at the same time as shown in Fig. 1a and b). (Note that throughout the paper the terms CO plume and sensitivity plume refer to the forward and backward model simulations, respectively.) Figure 1e and f shows the plan view and longitude-height cross section of the folded retroplume, derived from folding the mixing ratio and sensitivity fields shown in a–d, along with the contours showing the limits of the forward (blue) and backward (magenta) plumes from panels a–d. Note that throughout the paper we color any product derived from the forward model blue, from the backward model red and magenta, and from the folded retroplume green.

The folded retroplume clearly indicates the portions of the two plumes that successfully travel between the source and receptor. The folded retroplume also indicates the relative concentrations of the receptor-bound trace substance. A comparison of the folded retroplume with the forward CO and backward sensitivity contours shows that simply

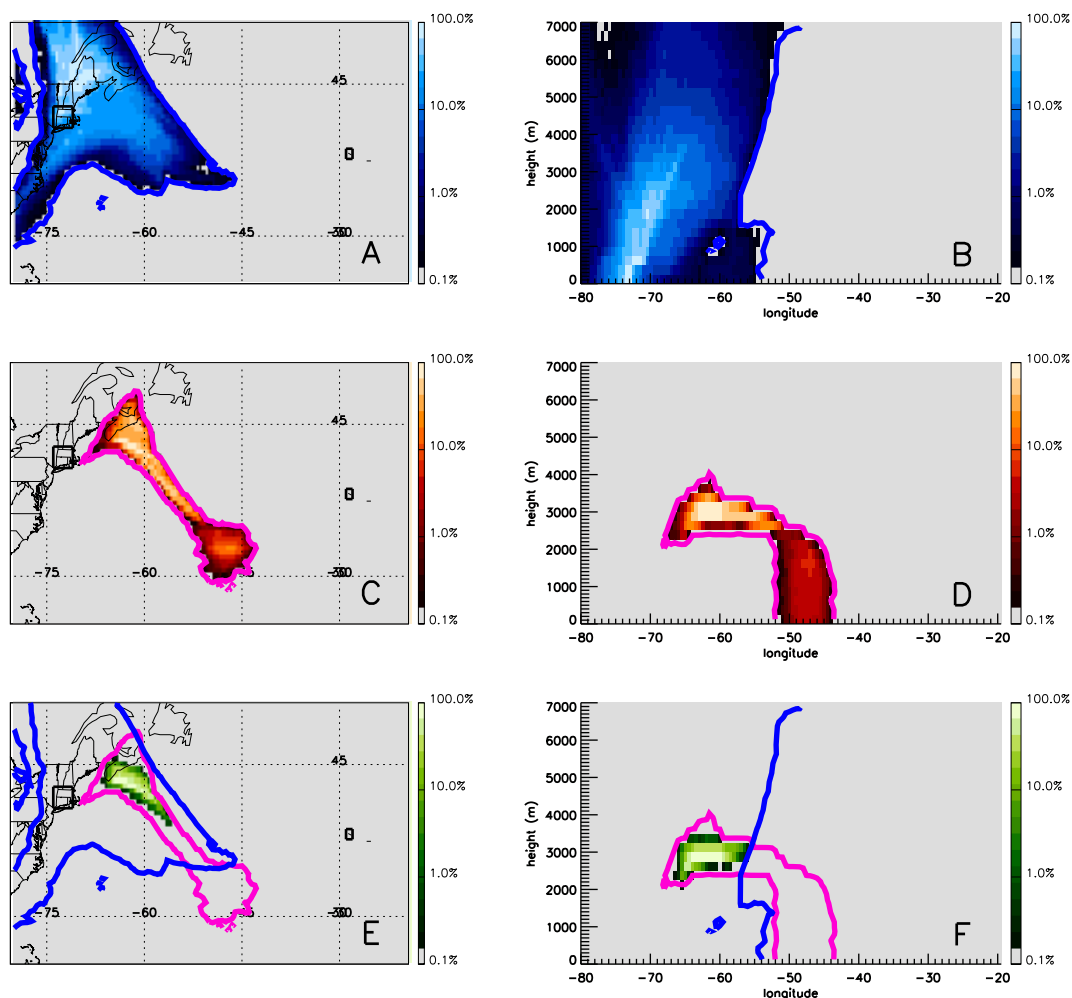


Fig. 1. The plan views (left column) and longitude-height cross sections (right column) of snapshots of the vertically integrated (left column) and horizontally integrated (right column) CO concentrations from a forward model simulation 1.5 days after release (**a** and **b**), the sensitivity field from a backward simulation 3 days upwind of the receptor (**c** and **d**), and the folded retroplume, or results of folding the concentration and sensitivity fields from (**a**)–(**d**) (**e** and **f**). The colors for the plumes are scaled according to the maximum value in each panel. Contours indicate the limits of the forward (blue) and backward (magenta) plumes and are drawn at 1% of the maximum value for each plot. The source and receptor boxes are outlined in black in (**a**), (**c**), and (**e**).

superimposing the forward and backward plumes can be misleading. In this case, the overlap of the two contour lines roughly define the folded retroplume in the vertical view (**f**). However, this is not the case for the plan view (**e**), as the folded retroplume only occupies a portion of the overlapping contours. This apparent inconsistency is the result of viewing vertically integrated fields. The trace substance plume and sensitivity field, while in the same vertical plane, are not actually colocated vertically. In Sect. 4, we provide an expanded comparison of the folded retroplume with the standard LPDM products in a sample analysis.

3 Method evaluation

3.1 Approach and methods for the detailed evaluation

The primary purpose of the evaluation is to determine how well the folded retroplume reconstructs the pathway of emissions from the source to the receptor. Here, we use particle trajectories from the LPDM model runs to evaluate the accuracy of the folded retroplume pathway. The secondary purpose of the evaluation is to examine the behavior of the UMRs along the transport pathway. As discussed above, the UMRs should be constant if no emissions are added to the forward model. Deviations in the UMRs, which indicate disagreement between the transport described by the forward

and backward models, will also be investigated using particle trajectories. Some degree of disagreement is expected, as a result of the random components in the models (turbulence and convection). Finally, we will relate the behavior of the UMRs to the accuracy of the folded retroplume pathway.

This evaluation used the same model simulations presented in Sect. 2.2, which focused on the transport of anthropogenic CO emissions from a source region near Boston, MA on the US east coast to a receptor cell located around the Pico Mountain observatory in the Azores Islands in the Central North Atlantic. The two are located approximately 3620 km from each other and the transport time from source to receptor was 4.4 days. This transport time and distance should be sufficient to allow for any deviations from the expected pathway and UMR to become apparent. While the source region was chosen arbitrarily, the timing of the transport scenario was selected by first running a forward simulation with continuous emissions, from April to June 2005. We then selected one of the periods with the largest CO mixing ratio in the receptor cell for further inspection, with no prior knowledge of the transport scenario.

In this evaluation, we will discuss two types of particles, termed positive particles and negative particles based on whether or not they travel from the source to receptor over the period analyzed. From the forward simulation, positive particles are the trace substance particles that arrive in the receptor cell during the release of the backward plume. From the backward simulation, positive particles are the sensitivity particles that arrive in the source cell at the release time for the forward simulation. Negative particles are all other particles, including particles that do not travel from the source to the receptor as well as particles that successfully travel between the source and receptor cells, but not within the time frame of interest.

Note that at most times and in most grid cells, there will be both positive and negative particles from both model directions. Since dispersion causes increasing separation over time between particles that are initially near one another, the amount of dispersion a plume has experienced affects the relative number of concurrent positive and negative particles. As the two plumes are tracked toward the receptor, the forward plume will disperse and the backward plume will coalesce. Thus, near the source and close to the release time, many negative forward particles should be located along with positive forward particles, as the forward trace substance plume has experienced relatively little dispersion. In contrast, near the receptor and release time, only a few negative forward particles should be collocated with the positive forward particles, as the forward trace substance plume should be highly dispersed. For the backward simulation, there should be few concurrent positive and negative backward particles near the source and many concurrent positive and negative particles near the receptor. Theoretically, no location should ever contain only negative particles from both model directions, nor should positive particles from one

model direction be located in a cell without positive particles from the other model direction. These two situations indicate differences in the transport described by the two model simulations. In practice, however, this can occur, due to the random model components, transport errors, or irreversible transport.

3.2 Detailed evaluation results

3.2.1 Detailed evaluation of the folded retroplume pathway

The time-integrated results from the forward and backward simulations used for the evaluation are shown in Fig. 2. Figure 2a, which shows the plan view of the vertically integrated CO concentration field, indicates that the bulk of the CO emissions are transported northward. These emissions move out of the plot window; later, however, some of these emissions travel southward, toward the receptor (present as the dark plume stretching south-east from the northern edge of the plot window). A significant portion of the CO plume also moves east and southeast, stretching from the US east coast, across the Atlantic, to the receptor. Figure 2b, which shows the time-height cross section of the horizontally integrated plots of the CO concentration field, indicates that the bulk of the CO is transported to higher altitudes during the first few days, though CO is distributed throughout all levels of the atmosphere during the last 3 days of transport. The plan view (Fig. 2c) and the time-height cross section (Fig. 2d) of the sensitivity plume indicate a number of pathways (i.e., areas of sensitivity) for air traveling to the receptor. The regions of highest sensitivity are in a fairly compact pathway starting from just off the east coast of Nova Scotia, where the air converged, coming equally from the North and the South (from the emissions region). There is a secondary region of sensitivity that also originates near the source region and travels over the Atlantic slightly farther south than the primary sensitivity region, converging with the primary transport pathway west-southwest of the receptor.

Near the receptor, the horizontal transport pathway of the emissions can be guessed from Fig. 2 by comparing the sensitivity with the CO concentrations, as there is only a small region of overlap of the sensitivity (Fig. 2c) and CO fields (Fig. 2a). However, near the source and in the intermediate transport, over the Atlantic, the pathway that emissions travel to the receptor is unclear from these plots alone. Both the CO and sensitivity occupy a large area, both horizontally and vertically. Thus, even for this simplified case, with only a puff of emissions into the forward model, determining the exact pathway (or pathways) taken by the emissions as they travel to the receptor is not possible from the plan view and cross sections plots alone. One would need to view snapshots (i.e., the distribution of the plume at a single time, as opposed to the time-integrated view shown in the figure) of the two plumes in order to do that. However, even when viewing

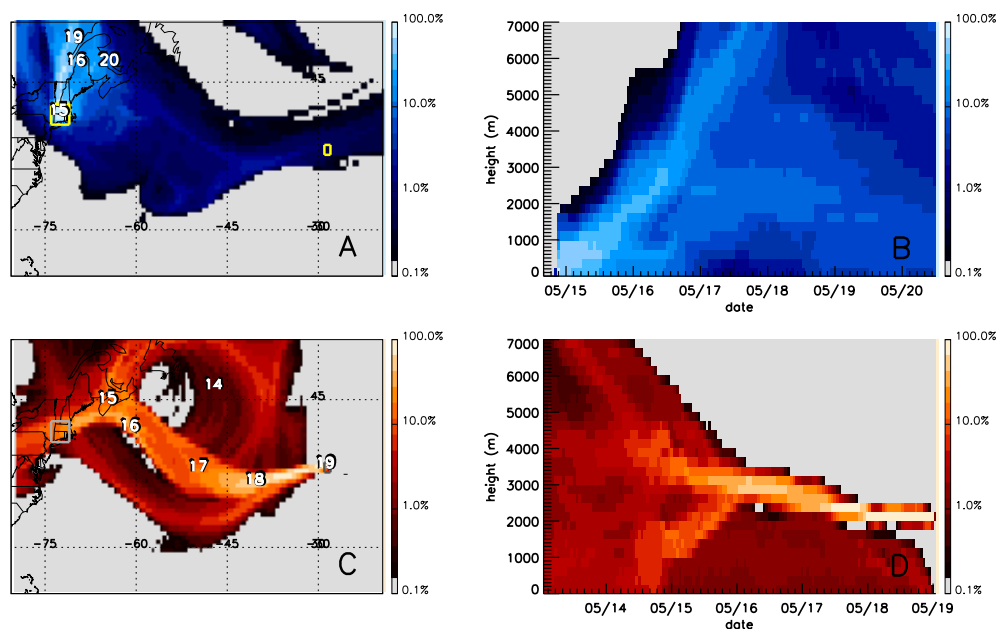


Fig. 2. The time integrated results from the forward and backward model simulations for the evaluation case. The plan view and time-height cross section of the CO plume from the forward simulation are shown in blue in (a) and (b), respectively. The plan view and time-height cross section of the sensitivity plume are shown in red in (c) and (d), respectively. The source and receptor volumes are outlined in yellow and gray in (a) and (c), respectively. The white numerals show the average location of the CO and sensitivity plumes at 00:00 UTC on the day of month indicated by the numbers.

multiple snapshots of the simulations, diagnosing the correct transport pathway can be difficult or impossible, as discussed in Sect. 2.2.

In contrast to the gridded output from the LPDM, the particle trajectories and the folded retroplume offer a clear view of the transport pathway between source and receptor. Figure 3 shows the plan view (a and c) and the time-height cross section (b and d) of the positive particles from the forward (a and b) and backward (c and d) model simulations. Figure 3e and f shows the plan view and the time-height cross section, respectively, of the folded retroplume pathway obtained from folding the two model simulations using Eq. (4).

In terms of the core transport pathway described by the three products, there was good overall agreement in both the horizontal and vertical pathways. All indicated lofting of emissions to 1–2 km in a daytime BL during the first few hours of transport. The emissions remained at this altitude, after the transition from a deeper continental daytime BL to a shallow nighttime marine BL left them located in the FT. Once in the FT, the emissions experienced slower ascent to about 3 km for approximately 1 day, where they remained for another day. Finally, during the last 2 days of transport, the emissions experienced a gradual descent from 3 km to the receptor at 2–2.25 km. The horizontal pathway shows that the emissions traveled northward along the coast to Nova Scotia, then traveled southeast before heading northeast again, toward the receptor. The common transport described by all

three products indicates that the folded retroplume successfully identifies the large-scale transport pathway between the source and receptor.

A comparison of the positive particle and folded retroplume pathways reveals three interesting features outside of the core transport pathway. Two of these features are regions where, due to the random components of the model, the pathways of the forward and backward positive particles differ. One such situation occurs during the initial day of transport (15–16 May) during the ascent from 1 km to 3 km (above the core transport pathway shown in panels b and d). The forward and backward maximum particle locations indicated in panel f encompass this region, indicating it is part of the source-to-receptor transport pathway. However, the smaller number of particle trajectories indicate that the probability of transport through this region is very low. The folded retroplume correctly identifies this low-probability region with a fairly small PMR, colored with darker greens and black. The second region where there is a difference between the forward and backward particle trajectories occurs during the last half of 16 May, when two forward positive particles stray below the core transport region (panel b). The very small number of trajectories from the forward model here indicates that this region is not part of the primary source-to-receptor transport pathway. The third feature is the presence of a few cells with a non-zero PMR that are entirely outside the limits of the positive particle pathway (i.e., a false positive PMR).

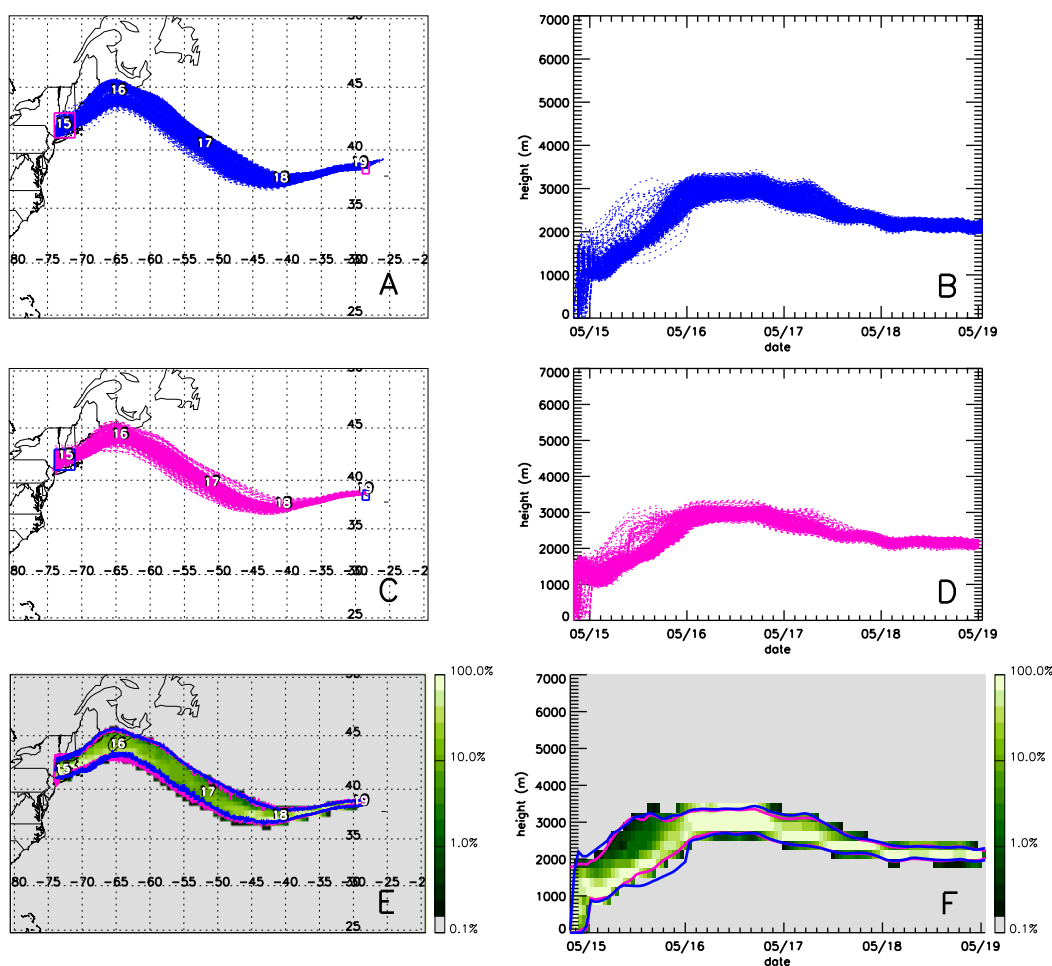


Fig. 3. The plan view and time-height cross section of the positive particle trajectories from the forward simulation are shown in blue in (a) and (b), respectively. The plan view and time-height cross section of the positive particle trajectories from the backward simulation are shown in magenta in (c) and (d), respectively. The plan view and time-height cross section of the folded retroplume are shown in green in (e) and (f), respectively. The solid blue and magenta lines in (e) indicate the location of the maximum and minimum positive particle latitudes at each longitude from the forward and backward model simulations, respectively. Similarly, the solid blue and magenta lines in (f) indicate the maximum and minimum positive particle altitudes at each day. The source and receptor volumes are outlined in magenta and blue in (a) and (c), respectively. The white numerals show the average location of the positive particles and the folded retroplume at 00:00 UTC on the day of month indicated by the numbers.

These cells are all adjacent to the positive particle pathway and contain only a small PMR (generally less than 1% of the maximum PMR and roughly 1% of the UMR). We show below that a blurring of the transport pathway up to one grid cell in size is the result of the use of gridded data. In summary, this evaluation indicates that the folded retroplume correctly identified the pathway of all but a small fraction (in this case, 3%) of the source-to-receptor pathway, with errors of up to about 1 grid cell in location.

3.2.2 Detailed evaluation of the UMRs

The UMRs are another important component of the folded retroplume as they can be used to estimate the timing and rate of emissions into or removal of trace substance from the

plume. The behavior of the UMRs can also be used to determine the degree of agreement in the transport as described by the forward and backward model. As discussed above, the UMR should be constant if no emissions are added to the model between time steps. Since we use a puff of emissions for this evaluation, the UMR should be constant at all times between release and arrival at the receptor. Thus, deviations from the expected UMR (i.e., a non-constant UMR) indicate when there are differences in the transport described by the forward and backward model simulations.

Figure 4 shows the UMRs at each upwind time for the folded retroplume. The left-hand axis indicates the absolute UMRs (pptv CO), and the right-hand axis indicates the UMRs normalized by dividing by the last UMR before arrival. (Any of the UMRs could have been used for




Fig. 4. Upwind mixing ratios (UMRs) at each transport time between departure at the source and arrival at the receptor for the evaluation simulation. The bottom axis indicates the date while the top axis indicates the upwind day number. The left-hand axis indicates absolute UMRs and the right-hand axis indicates the relative UMRs, normalized by the last UMR. Numbers and vertical dotted lines indicate periods discussed in detail in the text.

normalization, but we chose this value because there has been little deformation of the air mass between the sample time and the arrival time in the receptor due to the calm meteorological scenario. As a result, it should closely represent the UMR that would be calculated if the forward simulation could be sampled while in the receptor cell.)

There are clearly significant variations in the UMR values, approximately 35% in the negative direction and a little more than 45% in the positive direction. We consider two possible causes for these deviations. These deviations can be the result of irreversibility, which would indicate that the forward and backward simulations are not simulating the same transport. Irreversible transport can be caused by one or more of a number of factors, including the violation of the well-mixed criterion (Thompson, 1987), not maintaining a consistent representation of the mass of particles as air density changes (Lin et al., 2003), and errors induced by the interpolation of meteorology between grid points (Stohl, 1998). The two positive particles that stray from the primary transport pathway and do not coincide with positive backward particles may indicate that a small portion of this event is irreversible. However, the overall agreement between the forward and backward positive particle trajectories indicate that the primary transport pathway described for this event is reversible. Another source of these deviations could be the presence of sub-grid gradients in the CO or sensitivity field that are lost when gridded fields are calculated. In order to investigate this potential, we conducted detailed inspections of the CO concentrations, sensitivity, and PMR fields and the distributions of the positive and negative particles in the vicinity of the positive particles at several times, marked by the dotted vertical lines in Fig. 4.

This investigation determined that the low UMRs resulted from minor displacements (less than the size of a grid cell) between the groups of forward and backward positive particles, specifically in regions with a high mixing ratio or sensitivity gradient. For the transport scenario examined here, the forward CO plume near the receptor (time period 4 in Fig. 4) is a thin filament, on the order of 2–4 grid cells wide, and the positive forward particles are at the edge of this filament of CO (only the part of the CO plume that contains positive particles actually passes through the receptor cell). Thus, when the positive backward particles are displaced slightly from the positive forward particles, they are in a region with little CO, resulting in a very small PMR in those cells and a negative bias in the calculated UMR. A similar case can be found in period 3.

Around period 2, the roles of the sensitivity and CO plume begin to change. Around this time, the sensitivity plume splits (as it is followed backward in time), with one portion heading northeast and another portion (which contains the positive particles) heading southwest, towards the receptor. Meanwhile, the CO plume (as it is tracked forward in time) is also in the process of splitting in two. One portion is the filament that eventually travels to the receptor, and the other is the larger portion that travels northeast from the source. The positive forward and backward particles were still slightly displaced from one another. However, they were no longer located at the edge of their respective plumes, and thus no longer in a region of a high sensitivity or CO gradient. As a result of these conditions, the UMRs around this time are closer to the expected value.

Closer to the source region, however, a different situation results in UMRs with a positive bias. First, there is again a large CO gradient (as in the other periods). However, now there is a relatively large concentration of negative forward particles in these cells, because the plume has not dispersed much yet. Second, the sensitivity plume was more dispersed. The positive backward particles reside in more cells and are distributed more uniformly than the positive forward particles. As a result of these two factors, the positive backward particles are located in cells with a large number of negative forward particles. The sensitivity plume is therefore combined with significantly higher CO, resulting in a higher UMR. As the trace substance will be ubiquitous very near source regions in most model scenarios, this situation may occur frequently when UMRs are calculated close to the source region.

3.3 The impact of various model settings on the folded retroplume

Many model settings can affect how well transport is described, affecting the correlation between model simulations. Additionally, the way in which the output is saved can affect the number of positive and negative particles that are identified as being collocated. As the evaluation above

Table 1. Model settings used for evaluation.

Parameter name	Setting options
Averaging	On and off ¹
Temporal output interval	1 ¹ , 3, and 6 h
Spatial output grid ^a	0.5° ¹ , 1.0°, and 2.0°
Kernel	On and off ¹
Internal model time step ifine and ctl	5 and 20 ^{1,b}
Number of forward particles ^c	75 000 ^d and 500 000 ^{1,e}
Number of backward particles ^c	2500 ^f and 20 000 ^{1,g}

¹ Identified as preferred model settings for folded retroplume.

^a Used for both latitude and longitude simultaneously.

^b Used for both ifine and ctl simultaneously.

^c High and low number of particle pairs only run together.

^d Resulted in 20–40 positive particles.

^e Resulted in 250–625 positive particles.

^f Resulted in 25–35 positive particles.

^g Resulted in 250–650 positive particles.

demonstrated, these factors can in turn significantly affect the folded retroplume UMRs. We have evaluated the impact of the model settings listed in Table 1 upon the folded retroplume pathway and UMRs. We used the same release scenarios used in the evaluation, with every possible combination of the settings in Table 1 (144 simulations in total) in order to assess the impact of each model setting on the resulting folded retroplume pathway and UMRs. By running all possible groupings for these settings, we are able to evaluate the impact of changing one setting across all other possible settings. The model settings that were identified as the best settings (i.e., produced the most accurate pathway and UMRs) were used in the evaluation presented above. Here, we discuss the degree to which use of other model settings changed the evaluation results.

3.3.1 Folded retroplume pathway

Across all the model settings, the folded retroplume pathway was qualitatively similar to the results presented in the evaluation above. Larger output grid sizes naturally increased the size of the folded retroplume pathway, as the cells on the edge of the pathway were larger. The use of time-averaged output produced a mild ghosting effect, which is the superposition of negative particles that are located in the same cell but at different times during the averaging period. Thus, when averaged output was used, the folded retroplume pathway tended to be larger, with the occurrence of a few false positive PMRs along the edge of the transport pathway. Despite these two issues, the resultant folded retroplume pathway correctly identified the core transport pathway taken by the positive particle for all settings.

3.3.2 Folded retroplume UMRs

The general behavior of the UMRs were similar to those shown in Fig. 4: lower near the receptor, highly variable from 1 to 2.5 days upwind, relatively flat at approximately 3 days upwind, and very high at 4 days upwind, near the source. The higher positive bias in the UMR near the source region was present in all scenarios, indicating that no particular setting can help resolve this issue. This is not surprising, given the cause of this issue discussed in Sect. 3.2.2. The absolute value of the UMRs varied significantly with changes in the spatial size of the output grid, the frequency of output, and the use of average or instantaneous output, each of which we discuss further below. However, the other three model settings (the number of particles, the sampling kernel, and the model time steps, ifine and ctl) had little impact on the UMRs, and will not be discussed in detail. The number of particles were chosen so that forward and backward simulations both had roughly 30 or 600 positive particles for the small and large number of total particle sets, respectively. The lower number of particles was sufficient to return an accurate folded retroplume, which bodes well for future use of the method, as a lower number of particles can significantly decrease the computational time necessary for the forward simulations.

3.3.3 Spatial grid sizes

The size of the spatial grids can affect how positive and negative particles are associated with one another. A larger grid cell can either increase or decrease the UMR, depending on the circumstances. For example, consider a cell that extends vertically from the surface into the FT, in a case in which the pollution plume that would reach the receptor was traveling in the lower FT. This would result in all the positive forward and backward particles residing in only the top half of the cell. If this cell were over an emissions source, then the lower portion of the cell would have a large number of negative forward particles, released from the surface source. The use of this single cell would result in a significant overestimate of the PMR, since the forward particles in the BL would be included in the calculation. However, if this cell were not over an emissions source and the bottom half of the cell had no forward particles, the result would be an underestimate of the PMR in this cell, as the larger cell would dilute the mass of the positive forward particles over a larger region, giving a smaller mixing ratio, without affecting the sensitivity. In our analysis, the resulting UMRs either stayed the same or decreased by 5–10% with each increase in the grid cell size, though the decrease in the UMRs was more pronounced as the grid cell size was increased from 1° to 2°. Therefore, we recommend use of a grid size of 1° × 1°.

3.3.4 Averaged or instantaneous values and the length of averaging period

The folding method we propose depends on the collocation in both time and space of positive particles in order to successfully identify cells that contain the trace substance that will travel between the source and receptor. As noted in Sect. 3.3.1, time averaging produced a ghosting effect and produced false positive PMRs. This ghosting effect also produced a positive bias in PMRs within the folded retroplume pathway, when, over the course of the averaging period, an output grid cell spanned a region that included both the plume of interest (positive forward particles) and forward-model CO that did not ultimately reach the receptor (negative forward particles). This could occur, for example, when a sensitivity plume was located along the edge of the CO plume. The ghosting effect could cause the sensitivity plume to sample portions of the center of the CO plume, resulting in higher UMRs. Longer averaging periods can increase the likelihood that this ghosting effect can occur. At the shortest averaging period of 1 h, the averaged and instantaneous UMRs differed only slightly. The UMRs increased significantly as the length of the averaging period increased, with the UMRs roughly doubling at all times at each time interval increase. If UMR values are to be used quantitatively, we recommend instantaneous fields over averaged fields for a folded retroplume. Instantaneous values, however, increase the stochastic uncertainties of the output field. Since the 1 h averaged output and the instantaneous output were quite similar, a short averaging period (e.g., less than or equal to 1 h) may also be a good option, as the averaged fields will reduce the stochastic uncertainties without affecting the UMRs greatly.

3.4 Summary of evaluation

The transport pathway evaluation indicates that the folded retroplume does a good job of restoring the source-to-receptor transport pathway and that this pathway is quite robust over a variety of model settings. The UMR evaluation, however, indicates that large gradients in either the trace substance or sensitivity field combined with minor differences in transport can significantly impact the UMRs. We also found that the ubiquitous nature of the trace substance in the source region in combination with a well-dispersed sensitivity plume can lead to positive deviations in the UMRs. The use of larger grid sizes or an averaging period greater than 1 h can significantly degrade the accuracy of the UMRs, resulting in bias (mainly positive).

One important result from the evaluation is that significant deviations in the UMRs did not correlate with significant differences between the folded retroplume and positive particle transport pathways. Whether the UMRs were high or low, the correct cells were generally identified (i.e., the cells with positive particles), with the core transport and fringe cells appro-

priately corresponding to high- and low-probability transport regions. As a result, the folded retroplume pathway appears to be a robust product, even when differences in transport between the forward and backward model are indicated by variations in the UMRs.

4 Sample application

In this section, we present a sample analysis that contrasts the folded retroplume method with traditional methods using only standard gridded LPDM products. The analysis will serve to provide an example of the advantages of the folded retroplume method over traditional LPDM analysis methods. The sample analysis will again focus on the transport of US emissions to the Pico Mountain observatory, examining the transport scenario for an event observed at the Pico Mountain observatory from 21–23 April 2005. During the event, CO mixing ratios ranged from 120–180 ppbv, approximately 30–90 ppbv above the typical springtime background at the station, while FLEXPART indicated enhancements of 20–50 ppbv of CO. Ozone, nitrogen oxides, and non-methane hydrocarbons were also elevated during this period. The event is the second of two events discussed by Honrath et al. (2008).

As with the evaluation simulations, we use FLEXPART version 6.2, driven with ECMWF meteorological data, with North American CO emissions based on the EDGAR inventory (see Sect. 3.1 for more details). For the sample analysis, we chose settings that are fairly typical of FLEXPART applications, even though they deviated somewhat from the recommendations above, especially in terms of the averaging period used. CO emissions were released continuously over North America into the lowest 300 m of the atmosphere and carried in the model for 20 days, after which time they were dropped from the simulation. Particles for the backward simulation were released over a 1-h period centered on 21:00 UTC on 21 April, into a box bounded by 38–39° N latitude 29–28° W longitude, from an altitude of 2000–2500 m a.s.l. The output was saved on a 1° × 1° grid, with the top of the output levels at 0.3, 1, 1.5, 2, 2.5, 3, 4, 5, 7.5, 10, and 15 km. 6-h averages were saved and the sampling kernel was used. Particle trajectories from the backward model simulation were also saved and were used to confirm that the folded retroplume correctly captured the transport pathway of emissions between the source and receptor during the last 8 days of transport.

4.1 Comparison of the folded retroplume pathway with standard LPDM products

The sensitivity plume (Fig. 5a and c) shows fairly well organized transport originating over the US west coast about 7 days upwind. It also shows some sensitivity over a secondary region in the Central US, from Texas to the Great




Fig. 5. The time-height cross section (a and b) and plan views (c and d) of the sensitivity plume (red, a and c) and the folded retroplume (green, b and d) for the sample analysis. The plan-view color scales are based on the maximum specific volume weighted residence time ($3.1 \text{ s m}^3 \text{ kg}^{-1}$) and PMR (1.5 ppbv) for each plot. The white numerals show the average location of the sensitivity plume and folded retroplume at each upwind day, as indicated by the numbers. The time-height cross sections are normalized by the maximum value at each upwind time.

Lakes. Since the age of this secondary region is not apparent, snapshots of the sensitivity would be needed to determine its age. The time-height cross section shows that the sensitivity plume was in the lower levels of the atmosphere, where it would be able to pick up emissions if they were present, from about 6.5 to 4 days upwind. At 4 days upwind, the bulk of the sensitivity plume experienced relatively rapid uplift to the lower FT, where it stayed until arrival at the observatory. Thus, from these views of the sensitivity plume, one could conclude that the bulk of emissions would have been picked up during the leg of transport from the US west coast to Central, Northern US (i.e., between the 6 and 4 in Fig. 5a). A map of the source contributions (i.e., $\sum_j S_{j,t,(j',t')} \cdot E_{j,t}$ from Eq. 7, not shown) would show the true sources, primarily the Central US and only partly from the US west coast. However, tracking the emissions once away from the primary source region would be difficult due to the low sensitivity between this region and the compact transport pathway that occurs starting from approximately 2.5 days upwind.

In contrast to the sensitivity plume, the transport leg from the west coast almost disappears in the folded retroplume (Fig. 5b), indicating this transport pathway actually carries only a small amount of the trace substance. The folded retroplume shows that the emissions travel slowly northward in the lower atmosphere from upwind day 7 to 2. At 2 days upwind, all of the emissions were transported out of the BL into the lower FT, significantly later than indicated by the sensitivity plume. During the last 2 days of transport, the folded

retroplume and sensitivity plume indicate virtually identical transport, which is not surprising since the sensitivity indicates little dispersion in the last 2 days of transport.

In order to examine the causes for the differences between the standard and folded retroplume and to further demonstrate the utility of the folded retroplume, in Fig. 6 we show snapshots of the sensitivity plume (left column) and the folded retroplume (right column) plotted with contours of the total column CO from the forward simulation (blue lines). Snapshots are shown for 7(a and b), 5(c and d), 4(e and f), and 3(g and h) days upwind. For consistency among these plots, we have used the maximum sensitivity and PMR from all 4 plots for the color scale maximum.

We have selected contour levels and color scales for Fig. 6 that approximate typical usage for the CO and sensitivity plumes (e.g., Trickl et al., 2003), as some arbitrary selection of these settings was required. In some cases, it may be possible to adjust the contour levels and color scales such that features that were originally not apparent become visible. Often, however, this would result in confounding other areas of the figure. Additionally, adjusting these settings for each snapshot and product is not a realistic analysis approach and quite often would require the analyst to have a prior knowledge of the feature they are trying to identify (from, e.g., the folded retroplume or particle trajectories). Thus, it is reasonable to assume that the features identified by the folded retroplume but not the standard products, as presented here, would indeed be missed in most analyses of this type.




Fig. 6. The plan view of snapshots of the horizontally integrated sensitivity plume (red, left column) and the folded retroplume (green, right column) with forward CO contours (blue lines, left and right columns) at 7 (a and b), 5 (c and d), 4 (e and f), and 3 (g and h) days upwind. Sensitivity plume color scale based on the maximum specific volume weighted residence time for all 4 plots ($3.1 \text{ s m}^3 \text{ kg}^{-1}$). Folded retroplume color scale based on the maximum PMR for all 4 plots (1.5 ppbv). CO contours at 10, 20, 30, and 60% of the maximum column CO for all 8 plots (400 mg/m^2).




Fig. 7. The UMRs at each upwind time for the sample evaluation. The UMRs calculated using the appropriate CO age classes are shown in green, while the alternate UMRs calculated with the total CO are shown in black. The cumulative CO distributions from the forward model (blue boxes) and the backward model folded with emissions (magenta crosses) are also shown.

At 7 days upwind (Fig. 6a and b), there are significant differences between the sensitivity plume and the folded retroplume. The region of sensitivity over the Western US and Eastern Pacific (Fig. 6a) completely disappears in the folded retroplume (Fig. 6b), while the region of sensitivity over the Central US is emphasized, with relatively large PMRs. Considering the overlap of the sensitivity plume over the Central US with the CO contours, the lack of CO in the western lobe of the sensitivity field, and the altitude of the primary portion of the sensitivity plume (Fig. 5), this result is not surprising. However, it serves to demonstrate the cause for differences between the sensitivity plume and the folded retroplume. The emphasized area over the Central US with large PMRs, however, is considerably larger than the area indicated by the sensitivity plume and extends to regions with relatively low CO levels (i.e., it extends beyond the area enclosed by the CO contours).

At 5 days upwind (Fig. 6c and d), there remain significant differences between the sensitivity plume and the folded retroplume. The eastern lobe of the sensitivity plume has grown larger. The western lobe of the sensitivity plume now shows up in the folded retroplume, as it has picked up some emissions over the Western US. In the absence of selectively picking contours and color scales, it would be impossible to identify the western lobe as contributing to CO transported to the receptor, as the sensitivity plume and selected CO contours do not overlap.

By 4 days upwind (Fig. 6e and f), the primary eastern and secondary western lobes of the sensitivity plume and PMR fields have merged. A significant portion of the folded retroplume is now outside the CO contours. (That is, this feature would be missed using the standard products shown in panel e or by using superimposed contours of the sensitivity plume and forward output alone, like those shown in

Fig. 1. A minimum CO contour level of 1% of the maximum would be required to indicate that the sensitivity plume and CO fields overlap. However, if this level were used for all the plots, then the sensitivity plume and CO contours would completely overlap at all upwind times shown in Fig. 6, which would incorrectly indicate that all of the sensitivity plume carried emissions.)

By 3 days upwind (Fig. 6g and h), the two lobes have almost completely coalesced and the outline of the folded retroplume and sensitivity plume are fairly similar, though the colors indicate the distribution of the trace substance within the plume is still quite different than the distribution of the sensitivity field. The folded retroplume and sensitivity plume are virtually identical during the final 2 days of transport (shown in Fig. 5), which is expected, as the sensitivity plume was very narrow during this period.

4.2 Folded retroplume UMRs

The UMR distribution, shown in Fig. 7 (thick, green solid line marked with diamonds), can be used to identify details about the addition of emissions to the plume. Increasing UMRs indicate the time periods when emissions were being added to the plume. The slope of the UMR line during this period indicates the rate at which emissions were added. Finally, constant UMRs indicate periods when the sensitivity plume was not located over source regions. (Note that if removal processes were included, decreases in the UMR would identify times during which removal occurred, though the interpretation of changes in the UMR would be complicated by the competition between emission and removal. See Sect. 5.3 for a discussion on this topic.)

The distribution of the CO mixing ratio age classes in the receptor, which can be derived from the forward model (blue line in Fig. 7) and from the backward model (by folding the sensitivity plume with the emissions inventory, magenta line in Fig. 7), provide another means to determine the details about the timing and rate of the addition of emissions to the plume. Therefore, a comparison of these distributions with the UMRs can help evaluate how well the UMRs accomplish this task by differentiating between changes in the UMRs that result from differences in transport and those that result from emissions (and/or removal, were it used). Since the folded retroplume samples the forward simulation, the UMR at t days upwind should equal the sum of the CO mixing ratio age classes from the forward model that are greater than t days old:

$$\text{UMR}_t \approx \sum_{i=t}^{A_f} \chi_{t,(j'),i}, \quad (8)$$

where A_f is the number of days the trace substance is carried in the forward model, $\chi_{t,(j'),i}$ is the CO mixing ratio of age i (in days) in the receptor cell (j'), and the sum is over all available CO mixing ratio age classes in the cell.

The relatively small change in the UMRs and CO distributions from 0 ppbv to about 5 ppbv during 20 to 10 days upwind indicates that, during this time, the majority of the plume was not over an emissions region. From 10 to 4 days upwind, however, the UMRs (green) and sensitivity-derived CO (magenta) increased by about 17 ppbv and the forward CO increased by about 20 ppbv (about 62% and 72% of the final UMR, respectively). This increase indicates that this is the time period when the plume was over the emission region and CO was actively being added to the plume. Between 9 and 4 days upwind, the emission rate was fairly constant, at about 2 ppbv/day according to the UMRs. During the last 3 days of transport, however, CO distributions are relatively flat, indicating the plume was no longer in the emissions region, in agreement with the location of the folded retroplume shown in Figs. 5 and 6. The UMRs are also relatively flat between days 3 and 2, closely following the sensitivity-derived CO distribution. Between days 2 and 1, however, the UMRs increase significantly, coming closer to the forward CO distribution, and are flat again, following the CO from the forward simulation during the last day of transport.

In order to understand the differences between the UMRs and the forward CO mixing ratio distribution (23–24% from upwind day 7 to 2), we conducted an inspection of the CO plume, folded retroplume, sensitivity plume, and positive particles from the sensitivity plume during days 2–7 (not shown). This inspection revealed several features that, together, explain these low UMRs. First, the positive particles were not collocated with the areas with the maximum PMR, indicating minor differences in the transport between the forward and backward model simulations. Second, the gradient in the forward CO mixing ratio is fairly large over the area covered by the folded retroplume (varying by about 40% in the horizontal direction and about 85% in the vertical direction). The evaluation in Sect. 3 showed that in situations like this, with a large gradient in either the sensitivity or trace substance plume, the UMRs can be sensitive to minor displacements in the positive particle locations.

The magnitude of the deviations of the UMR can be put into perspective by considering the magnitude of the disagreement between the forward CO mixing ratio age class distribution (blue line in Fig. 7) and distribution derived from folding the sensitivity with the emissions inventory (magenta line in Fig. 7). The differences between the forward (blue) and backward (magenta) derived CO distributions is partially the result of the differences in transport described by the two model simulations (and partially the result of differences in the number of particles used). Seibert and Frank (2004) showed that even with a short travel distance, the difference between the two can be in excess of 10%. The difference between the CO derived from the forward output and backward output convolved with the emissions inventory is not abnormal for FLEXPART, based on our own comparison of the two. (Our analysis consisted of a comparison of the total CO mixing ratio obtained from the forward model to the CO

mixing ratio obtained from simulations of backward model, initiated at the Pico Mountain observatory, convolved with the emissions inventory. The comparison included approximately 6000 CO mixing ratios, covering 2 years worth of data.) The UMRs are in fact bounded by the forward and backward CO distributions, which are derived from established modeling methods. This suggests that the deviations in the UMR during these times are reasonable.

Archiving the full spectrum of forward trace substance age classes can require a significant amount of storage space. For the example presented here, we saved CO age classes at a 6-h resolution for the whole Northern Hemisphere in order to appropriately match sensitivity and CO mixing ratios. (This one-month simulation required 15.5 GB of storage.) As a result, it is preferable for storage reasons to save only the total CO mixing ratio. To evaluate the effect that this would have on the UMR results, we have included the UMRs resulting from PMRs calculated by folding the sensitivity plume with all available CO mixing ratio age classes. Thus, Eq. (5) becomes

$$\text{PMR}_{j,t} = S_{j,t,(j',t')} \cdot \sum_{i=0}^{A_f} \chi_{j,t,i} \quad (9)$$

These UMRs are shown in Fig. 7 (black line marked with asterisks). They match the other UMRs quite well during the last 10 days of transport. However, from upwind days 10 to 20, there is an obvious positive bias in these UMRs, resulting from the sensitivity being folded with CO that will be dropped from the forward model before it can be transported to the receptor (i.e., the particles will reach the maximum age allowed and will thus be dropped from the forward simulation). Since we are mainly interested in the last ~10 days of transport, when the bulk of the emissions have been added to the plume and transported to the receptor, this would not be an issue with the example presented here. Deviations of the UMRs of this kind are most likely to occur at the greatest upwind times, when very old CO emissions can be folded with sensitivity many days upwind of the receptor.

5 Alternate backward LPDM combinations

The method we have presented uses the transport sensitivity from the backward formulation of a model to estimate the location and distribution of the trace substance in the forward model that will eventually arrive at the receptor. This is only one of many potential applications that combine the transport sensitivity from a backward LPDM with gridded fields from a number of sources. The simplest example, given by Seibert and Frank (2004), is the convolution of the sensitivity plume with an emissions inventory to determine the age class distribution of the mixing ratio of the trace substance in the receptor cell. Here, we discuss a few possibilities that focus on the calculation of UMRs and similar estimates of the physical attributes of the sensitivity plume at upwind times.

5.1 Chemical transport models

UMRs calculated from the folding of a chemical transport model (CTM) with the backward LPDM would allow for an evaluation of the mixing ratios in the portions of the CTM sampled by the sensitivity plume. This would give an estimate of the chemical transformations occurring in the CTM from the Lagrangian perspective provided by the sensitivity plume and would provide insight into the chemical and transport processes in the CTM. A similar application, in which the adjoint of transport was used to sample a CTM was given by Vukićević and Hess (2000) and Hess and Vukićević (2003). Vukićević and Hess (2000) determined the adjoint of the CTM HANK to find the transport sensitivity of a non-reactive trace substance. These sensitivity fields were equivalent to those obtained from a backward LPDM. Hess and Vukićević (2003) then folded the adjoint sensitivity fields with output fields from the forward mode of the CTM to compute UMRs (or, as described by Hess and Vukićević (2003), the sensitivity-weighted average of these fields).

In order to calculate UMRs from a CTM, ideally, the adjoint of transport in the CTM would be used because transport would be described the same in both model directions. However, the backward LPDM is also an attractive alternative, particularly when the adjoint of the CTM is not already available. The UMRs calculated in this manner will equal the mixing ratio in the receptor if the LPDM and CTM transport are sufficiently similar, no emissions are added, no chemistry occurs, and, if a passive trace substance is used in the LPDM, no removal has occurred within the plume. The agreement of transport between the CTM and LPDM could be tested by comparing forward model simulations using a passive trace substance and identical emissions sources. Changes in UMRs over time can thus be used to evaluate the net impact of emissions, removal and chemistry in the CTM from a Lagrangian perspective.

5.2 Meteorological fields, measurement and satellite data, and other applications

Other potential applications exist, since any data that are available on a 3-D grid can be combined with the backward LPDM to calculate sensitivity-weighted fields. Such 3-D data include meteorological fields, high-resolution measurements, and satellite data. Folding meteorological fields with the backward LPDM sensitivity provides the sensitivity-weighted upwind meteorological conditions for the sensitivity plume (e.g., the average temperature of the plume at each upwind time). This could be useful for determining the physical conditions of the air that will ultimately reach the receptor, which can be used to drive chemical reactions in a Lagrangian box model (e.g., Evans et al., 2000).

UMRs calculated from atmospheric measurements provide an additional estimation of the chemical and physical properties of a plume upwind of a receptor. Fields of at-

mospheric composition derived from satellite observations or from intensive in-situ aircraft sampling could be analyzed by folding with the backward LPDM results to obtain UMRs. In these cases, the evolution of the air mass could be examined as it traveled to the receptor by comparing the UMRs over time.

5.3 Wet and dry removal

The discussion of the method has thus far focused only on simulations that do not include any removal processes (i.e., passive trace substances). The inclusion of removal processes, however, has important and interesting implications but requires careful consideration. Generally, removal should either be included or excluded in both of the paired model simulations. By comparing the UMRs from two sets of simulations – one from two simulations with no removal (subscript *nr* – UMR_{*nr*}) and one from two simulations with removal (subscript *wr* – UMR_{*wr*}), the timing of removal can be determined. For example, consider UMRs calculated at three time periods, one close to the source, t_s , one near the receptor, t_r , and an intermediate time, t_i . If no emissions are added to the air mass from t_s to t_r , then UMR_{*nr*} will be constant at all three times. However, if removal occurs between either t_s and t_i or t_i and t_r , then UMR_{*wr*} will decrease after the removal has occurred. Therefore, a comparison of the two UMRs will indicate the timing of removal. Once the timing has been determined from this method, the PMR_{*wr*}s can be compared with the PMR_{*nr*}s to identify where this removal occurred.

6 Practical considerations

There are a number of practical considerations relating to the model settings and parameters that should be taken into account when creating folded retroplumes. Two of the most important of these are discussed in this section and are expected to be of the greatest interest to readers wanting to conduct their own calculations of this type. This discussion is limited to the LPDM FLEXPART, but may be applied to other similar models by analogy.

6.1 FLEXPART output time stamp

In FLEXPART, the output is saved at the end of the averaging period. The forward and backward model simulations, however, are integrated in opposite directions in time. Thus, if averaged (normal) output is used, in order to appropriately match the model output, the length of the output interval must be added to the backward model output times, before identical forward and backward times are selected for folding.

6.2 Output unit options

The derivation presented in Sect. 2.1 used the volume mixing ratio and a form of the sensitivity that also has units of mixing ratio. Neither of these units are native to FLEXPART. Here, we discuss native FLEXPART output units that can be paired, along with some conversion factors, to obtain UMR units of volume mixing ratio. One such pairing is the mass missing ratio for the forward mode (with model settings `ind_source=1` and `ind_receptor=2`) and the residence time from the backward mode (`ind_source=1` and `ind_receptor=1`). In this case, the conversion factor necessary to obtain volume mixing ratios are the ratio of the mass of air to the mass of the trace substance (which serves to convert the mass mixing ratio to the volume mixing ratio) and the inverse of output interval (which serves to convert the residence time to the mixing ratio sensitivity). (Note, the inverse of the output interval may or may not be necessary with instantaneous output, depending on how the modifications are made to the model code.) Another pair of native FLEXPART output units is concentration from the forward mode (`ind_source=1` and `ind_receptor=1`) and the specific volume weighted residence time (SVWRT, `ind_source=1` and `ind_receptor=2`). The conversion factors are identical to the first pair, the ratio of the mass of air to the mass of the trace substance and the the inverse of output interval. Seibert and Frank (2004) and Stohl et al. (2005) review other potential model outputs.

7 Summary and conclusions

This paper introduced a new method, the product of which we call a folded retroplume, that identifies source-to-receptor transport pathways using standard gridded products from a Lagrangian particle dispersion model (LPDM). The folded retroplume can be used to determine the transport pathway of only those emissions that arrive at the receptor at a designated time and to estimate the timing and location of emission and removal processes within the model.

An evaluation was conducted to determine the ability of the folded retroplume to identify the source-to-receptor transport pathway and to consistently calculate the expected sensitivity-weighted mixing ratios (upwind mixing ratios or UMRs) along the transport pathway. A comparison of the folded retroplume pathway with particle trajectories from both the forward and backward LPDM simulations showed that the folded retroplume was consistently able to reproduce the source-to-receptor transport pathway across a wide variety of model settings. Minor differences between the folded retroplume and particle pathways along the edge of the core transport pathway were found, but were limited to one or two grid cells (typically 10–100 km), a fairly minor difference when compared to the length of the transport pathway analyzed (3620 km). The UMRs may be expected to be constant at all times between the source and receptor in

the evaluation simulations. However, significant variations in UMRs were found to occur across all model settings. The best model settings produced UMRs that typically deviated from the expected value by 20–40%, while other model settings produced deviations exceeding 100%. These deviations resulted primarily from errors induced by using gridded data when large gradients in the trace substance and sensitivity fields were present and as well as minor differences between the source-to-receptor transport described by the forward and backward model simulations. The test simulations, however, provided a rather rigorous test of the folded retroplume properties, particularly the UMRs, as the evaluation transport scenario produced high gradients.

A sample analysis of the transport of North American CO to a monitoring station located in the Azores Islands contrasted the folded retroplume with traditional LPDM analysis tools. The folded retroplume made it possible to identify transport features that were difficult or impossible to determine with the traditional LPDM analysis techniques. The UMRs proved to be useful as a tool to determine the timing and rate at which emissions were added to the plume. Had the sample analysis included removal processes, the timing and rate of removal could also have been determined. We conclude that the folded retroplume is better than traditional LPDM analysis techniques for determining source-to-receptor pathways because it is significantly easier and more accurate than the alternative (separately viewing all the components of the LPDM). The folded retroplume also provides information unavailable from traditional LPDM products, such as the location and timing of removal processes and the spatial distribution of the trace substance between the source and receptor.

The large deviations of the UMRs in the evaluation raised the issue of the reversibility of the transport scenario. Even though it was ultimately determined that the transport scenario was reversible, the issue of the reversibility of any particular transport scenario is important. While large variations of the UMR can occur under certain conditions even in the absence of irreversibility (as discussed in Sect. 3.2.2), UMRs that behave as expected indicate that the source-to-receptor transport pathway is correct and that the transport simulation was reversible. The example transport scenario presented in Sect. 4 serves to illustrate this idea. The UMRs in the example were within the bounds of the corresponding equivalent mixing ratios calculated from the forward and backward simulations, and the folded retroplume pathway was found to agree well with the positive particle trajectories from the backward simulation. Thus, UMRs may provide a means to determine when to accept the source-to-receptor transport pathway indicated by the folded retroplume as correct. In this respect, the folded retroplume may be superior to both a backward simulation and particle trajectories, neither of which offers a means to determine reversibility of the transport scenario.

The folded retroplume should be most useful for the analysis of measurements made in relatively remote regions of the atmosphere, where pollution impacts result from emissions that have traveled long distances. In these cases, emissions are often significantly dispersed, and the source-to-receptor pathway may be difficult to discern from standard LPDM products. This scenario is relatively typical of many monitoring stations (e.g., Carpenter et al., 1997 and Kentarchos et al., 2000) and aircraft-based sampling efforts (e.g., Cooper et al., 2001 and Lewis et al., 2007).

The LPDM folding technique can also be applied using the backward LPDM to sample fields of mixing ratios, concentrations, and meteorological data from other sources in order to explore gridded data fields in a Lagrangian framework. When paired with the concentration output from a CTM, the method could be used to probe the chemistry occurring in the CTM within the confines of the sensitivity plume. Applied to satellite or field measurements, the method could be used to examine the chemical and physical evolution of a plume in the atmosphere. Meteorological conditions extracted with the method could be used to understand the conditions experienced by an air mass during transport and can be used in a trajectory chemical transport model or box model. These alternate applications of the technique offer new opportunities to study physical and chemical transformations that occur in the atmosphere.

Acknowledgements. We thank D. Henriques, Institute of Meteorology, Ponta Delgada, Portugal, for providing the ECMWF data used in this work and Andreas Stohl, Norsk Institutt for Luftforskning (NILU), Kjeller, Norway for providing the FLEXPART model and assisting in its use. We also wish to thank the reviewers, especially P. Seibert, for their suggestions for improving the manuscript. This research was supported by grants from the National Science Foundation (ATM-0535486 and ATM-0720955).

Edited by: A. Stohl

References

- Carpenter, L. J., Monks, P. S., Bandy, B. J., Penkett, S. A., Galbally, I. E., and Meyer, C. P.: A study of peroxy radicals and ozone photochemistry at coastal sites in the northern and southern hemispheres, *J. Geophys. Res.*, 102, 25417–25427, 1997.
- Cooper, O., Moody, J., Parrish, D., Trainer, M., Ryerson, T., Holloway, J., Hübler, G., Fehsenfeld, F., Oltmans, S., and Evans, M.: Trace gas signatures of the airstreams within North Atlantic cyclones: Case studies from the North American Regional Experiment (NARE '97) aircraft intensive, *J. Geophys. Res.*, 106, 5437–5456, 2001.
- Cooper, O., Stohl, A., Eckhardt, S., Parrish, D., Oltmans, S., Johnson, B., Nédélec, P., Schmidlin, F., Newchurch, M., Kondo, Y., and Kita, K.: A springtime comparison of tropospheric ozone and transport pathways on the east and west coasts of the United States, *J. Geophys. Res.*, 110, D05S90, doi:10.1029/2004JD005183, 2005.
- ECMWF: Users guide to ECMWF products 4.0, Tech. Rep. Meteorological Bulletin M3.2, European Center for Medium-Range Weather Forecasts (ECMWF), Reading, UK, 2005.
- Errico, R. M.: What is an adjoint model?, *B. Am. Meteorol. Soc.*, 78, 2577–2591, 1997.
- Evans, M., Shallcross, D., Law, K., Wild, J., Simmonds, P., Spain, T., Berrisford, P., Methven, J., Lewis, A., McQuaid, J., Pilling, M., Bandy, B., Penkett, S., and Pyle, J.: Evaluation of a Lagrangian box model using field measurements from EASE (Eastern Atlantic Summer Experiment) 1996, *Atmos. Environ.*, 34, 3843–3863, 2000.
- Flesch, T., Wilson, J., and Yee, E.: Backward-Time Lagrangian Stochastic Dispersion Models and Their Application to Estimate Gaseous Emissions, *J. Appl. Meteorol.*, 34, 1230–1332, 1995.
- Han, Y.-J., Holsen, T. M., Hopke, P. K., and Yi, S.-M.: Comparison of back-trajectory based modeling and Lagrangian backward dispersion modeling for location sources of reactive gaseous mercury, *Environ. Sci. Technol.*, 23, 1715–1723, 2005.
- Hess, P. and Vukićević, T.: Intercontinental transport, chemical transformations, and baroclinic systems, *J. Geophys. Res.*, 108(D12), 4354, doi:10.1029/2002JD002798, 2003.
- Honrath, R., Helmig, D., Owen, R., Parrish, D., and Tanner, D.: Non-methane hydrocarbons (NMHC) at Pico Mountain, Azores: 2. Event-specific analysis of the impacts of mixing and photochemistry on hydrocarbon ratios, *J. Geophys. Res.*, 113, D20S92, doi:10.1029/2008JD009832, 2008.
- Kentarchos, A. S., Roelofs, G. J., Lelieveld, J., and Cuevas, E.: On the origin of elevated surface ozone concentrations at Izana Observatory, Tenerife during late March 1996, *Geophys. Res. Lett.*, 27, 3699–3702, 2000.
- Lehmann, R., von der Gathen, P., Rex, M., and Streibel, M.: Statistical analysis of the precision of the Match method, *Atmos. Chem. Phys.*, 5, 2713–2727, 2005, <http://www.atmos-chem-phys.net/5/2713/2005/>.
- Lewis, A. C., Evans, M. J., Methven, J., Watson, N., Lee, J. D., Hopkins, J. R., Purvis, R. M., Arnold, S. R., McQuaid, J. B., Whalley, L. K., Pilling, M. J., Heard, D. E., Monks, P. S., Parker, A. E., Reeves, C. E., Oram, D. E., Mills, G., Bandy, B. J., Stewart, D., Coe, H., Williams, P., and Crosier, J.: Chemical composition observed over the mid-Atlantic and the detection of pollution signatures far from source regions, *J. Geophys. Res.*, 112, D10S39, doi:10.1029/2006JD007584, 2007.
- Lin, J., Gerbig, C., Wofsy, S., Andrews, A., Daube, B., David, K., and Grainger, C.: A near-field tool for simulating the upstream influence of atmospheric observations: The Stochastic Time-Inverted Lagrangian Transport (STILT) model, *J. Geophys. Res.*, 108(D16), 4493, doi:10.1029/2002JD003161, 2003.
- Methven, J., Arnold, S., O'Connor, F., Barjat, H., Dewey, K., Kent, J., and Brough, N.: Estimating photochemically produced ozone throughout a domain using flight data and a Lagrangian model, *J. Geophys. Res.*, 108(D9), 4271, doi:10.1029/2002JD002955, 2003.
- Methven, J., Arnold, S., Stohl, A., Evans, M., Avery, M., Law, K., Lewis, A., Monks, P., Parrish, D., Reeves, C., Schlager, H., Atlas, E., Blake, D., Coe, H., Crosier, J., Flocke, F., Holloway, J., Hopkins, J., McQuaid, J., Purvis, R., Rappenglück, B., Singh, H., Watson, N., Whalley, L., and Williams, P.: Establishing Lagrangian connections between observations within air masses crossing the Atlantic during the Interna-

- tional Consortium for Atmospheric Research on Transport and Transformation experiment, *J. Geophys. Res.*, 111, D23S62, doi:10.1029/2006JD007540, 2006.
- Neuman, A., Parrish, D., Trainer, M., Ryerson, T., Holloway, J., Nowak, J., Swanson, A., Flocke, F., Roberts, J., Brown, S., Stark, H., Sommariva, R., Stohl, A., Peltier, R., Weber, R., Wollny, A., Sueper, D. T., Hubler, G., and Fehsenfeld, F.: Reactive nitrogen transport and photochemistry in urban plumes over the North Atlantic Ocean, *J. Geophys. Res.*, 111, D23S54, doi:10.1029/2006JD007010, 2006.
- Olivier, J., Bouwman, A., der Maas, C. V., Berdowski, J., Veldt, C., Bloos, J., Visschedijk, A., Zandveld, P., and Haverlag, J.: Description of EDGAR Version 2.0: A set of global emission inventories of greenhouse gases and ozone-depleting substances for all anthropogenic and most natural sources on a per country basis and on $1^\circ \times 1^\circ$ grid, Tech. Rep. 771060 002 and R96/119, National Institute of Public Health and the Environment (RIVM) and Netherlands Organization for Applied Scientific Research (TNO), 1996.
- Owen, R. C., Cooper, O. R., Stohl, A., and Honrath, R. E.: An analysis of the mechanisms of North American pollutant transport to the central North Atlantic lower free troposphere, *J. Geophys. Res.*, 111, D23S58, doi:10.1029/2006JD007062, 2006.
- Rex, M., von der Gathen, P., Harris, N., Lucic, D., Knudsen, B., Braathen, G., Reid, S., Backer, H. D., Claude, H., Fabian, R., Fast, H., Gil, M., Kyro, E., Mikkelsen, I., Rummukainen, M., Smit, H., Stahelin, J., Varotsos, C., and Zaitcev, I.: In situ measurements of stratospheric ozone depletion rates in the Arctic winter 1991/1992: A Lagrangian approach, *J. Geophys. Res.*, 103, 5843–5853, 1998.
- Seibert, P. and Frank, A.: Source-receptor matrix calculation with a Lagrangian particle dispersion model in backward mode, *Atmos. Chem. Phys.*, 4, 51–63, 2004, <http://www.atmos-chem-phys.net/4/51/2004/>.
- Stohl, A.: Computation, accuracy and applications of trajectories – a review and bibliography, *Atmos. Environ.*, 32, 947–966, 1998.
- Stohl, A. and Trickl, T.: A textbook example of long-range transport: simultaneous observation of ozone maxima of stratospheric and North American origin in the free troposphere over Europe, *J. Geophys. Res.*, 104, 30445–30462, 1999.
- Stohl, A., Wotawa, G., Seibert, P., and Kromp-Klob, H.: Interpolation errors in wind fields as a function of spatial and temporal resolution and their impact on different types of kinematic trajectories, *J. Appl. Meteorol.*, 34, 2149–2165, 1995.
- Stohl, A., Eckhardt, S., Foster, C., James, P., Spichtinger, N., and Seibert, P.: A replacement for simple back trajectory calculations in the interpretation of atmospheric trace substance measurements, *Atmos. Environ.*, 36, 4635–4648, 2002a.
- Stohl, A., Trainer, M., Ryerson, T. B., Holloway, J. S., and Parrish, D. D.: Export of NO_y from the North American boundary layer during 1996 and 1997 North Atlantic Regional Experiments, *J. Geophys. Res.*, 107(D11), 4131, doi:10.1029/2001JD000519, 2002b.
- Stohl, A., Forster, C., Eckhardt, S., Spichtinger, N., Huntrieser, H., Heland, J., Schlager, H., Wilhelm, S., Arnold, F., and Cooper, O.: A backward modeling study of intercontinental pollution transport using aircraft measurements, *J. Geophys. Res.*, 108(D12), 4370, doi:10.1029/2002JD002862, 2003.
- Stohl, A., Cooper, O. R., Damoah, R., Fehsenfeld, F. C., Forster, C., Hsie, E.-Y., Hübler, G., Parrish, D. D., and Trainer, M.: Forecasting for a Lagrangian aircraft campaign, *Atmos. Chem. Phys.*, 4, 1113–1124, 2004, <http://www.atmos-chem-phys.net/4/1113/2004/>.
- Stohl, A., Forster, C., Frank, A., Seibert, P., and Wotawa, G.: Technical note: The Lagrangian particle dispersion model FLEX-PART version 6.2, *Atmos. Chem. Phys.*, 5, 2461–2474, 2005, <http://www.atmos-chem-phys.net/5/2461/2005/>.
- Thompson, A.: Criteria for the selection of stochastic models of particle trajectories in turbulent flows, *J. Fluid Mech.*, 180, 529–556, 1987.
- Trickl, T., Cooper, O. R., Eisele, H., James, P., Mücke, R., and Stohl, A.: Intercontinental transport and its influence on the ozone concentrations over Europe: Three case studies, *J. Geophys. Res.*, 108, 8530, doi:10.1029/2002JD002735, 2003.
- Vukićević, T. and Hess, P.: Analysis of tropospheric transport in the Pacific Basin using the adjoint technique, *J. Geophys. Res.*, 105, 7213–7230, 2000.