



Local fractions – a method for the calculation of local source contributions to air pollution, illustrated by examples using the EMEP MSC-W model (rv4_33)

Peter Wind^{1,2}, Bruce Rolstad Denby¹, and Michael Gauss¹

¹Climate Modelling and Air Pollution Division, Research and Development Department, Norwegian Meteorological Institute (MET Norway), P.O. 43 Blindern, 0313 Oslo, Norway

²Department of Chemistry, UiT – The Arctic University of Norway, 9037 Tromsø, Norway

Correspondence: Peter Wind (peter.wind@met.no)

Received: 17 October 2019 – Discussion started: 11 November 2019

Revised: 14 February 2020 – Accepted: 26 February 2020 – Published: 30 March 2020

Abstract. We present a computationally inexpensive method for individually quantifying the contributions from different sources to local air pollution. It can explicitly distinguish between regional–background and local–urban air pollution, allowing for fully consistent downscaling schemes.

The method can be implemented in existing Eulerian chemical transport models and can be used to distinguish the contribution of a large number of emission sources to air pollution in every receptor grid cell within one single model simulation and thus to provide detailed maps of the origin of the pollutants. Hence, it can be used for time-critical operational services by providing scientific information as input for local policy decisions on air pollution abatement. The main limitation in its current version is that nonlinear chemical processes are not accounted for and only primary pollutants can be addressed.

In this paper we provide a technical description of the method and discuss various applications for scientific and policy purposes.

1 Introduction

The origin of atmospheric pollutants within a given region is one of the fundamental questions of air quality research. Degradation of air quality, either temporary or sustained, is often the result of both local and long-range-transported air pollution, originating from anthropogenic but also natural emission sources. Anthropogenic emissions are due to a large

number of different categories such as road traffic, industrial point sources and large area sources.

In order to devise optimal strategies for air pollution abatement, for example short-term or long-term emission reduction measures, air quality managers need to have access to reliable scientific knowledge about the origin of air pollution. Typical questions include the following: (a) by what amount can local air pollution be reduced through local measures only, and in which cases will regional or countrywide measures be necessary? (b) What will be the benefit of emission reduction measures imposed on one or several specific emission sectors? (c) Will these measures be efficient on a short time frame or should they be implemented on a longer-term basis?

Many different methods exist to extract information about the origin of air pollution (e.g., Thunis et al., 2019; Clappier et al., 2017). Some of them are based on measurements of the chemical composition of air masses in the region or interest (receptor region). Such a “chemical fingerprint” can then give hints on the origin (source region or sector) of the pollutants. Most methods, however, are based on models as these can be readily applied to scenario calculations as well. Chemical transport models (CTMs), in particular, are efficient mathematical tools that treat the emission sources, transport, chemical conversion and loss mechanisms of air pollutants in a consistent way and allow different scenarios to be assessed within a reasonable amount of computing time.

The simplest method to evaluate the importance of different emission sources in a CTM is the “direct” method (e.g., Folberth et al., 2012), sometimes also referred to as an “anni-

hilation” or “brute force” method, whereby the same model simulation is repeated with and without including a chosen emission source. The difference in pollutant concentrations in the receptor region can then be attributed to the chosen emission source (impact). In order to stay within quasi-linearity one can choose to reduce the emission source by only a small amount. This is usually referred to as a “perturbation method” (e.g., Jacob et al., 1999; Fiore et al., 2009) and is well suited to simulate the effects of policy measures to reduce emissions from certain sectors by a certain amount. However, one of the drawbacks of this method is that for each source contribution a new, independent simulation must be performed.

Many chemical processes in the atmosphere are nonlinear. For example, a doubling of the emissions from one specific source will not necessarily double its contribution to air pollution levels. This also implies that the sum of contributions (from individual sources) calculated by the direct method (or by perturbation methods) will in general not be equal to the total air pollution level calculated in a simulation in which emissions from all sources are included in full. Consequently, one has to distinguish between two different questions: (1) what is the effect of a change in emissions from individual sources on air pollution (air pollution sensitivity)? (2) What are the contributions of individual sources to air pollution (source apportionment)? Due to nonlinearities, question 2 cannot be answered by reducing the emissions of individual sources to zero one by one. An alternative approach to estimate contributions from individual sources in model calculations is a technique known as “tagging”, which distinguishes chemically identical molecules according to their sources. In the calculation the molecule is labeled (e.g., by a separate index) according to its source and then keeps this label during transport and chemical transformations. When analyzing air pollution levels within a given receptor area, the fractions of molecules with different labels can be considered separately, thereby giving an estimate of the contributions from the different sources. A series of methods have been proposed to address the contribution from different sources based on the tagging method (e.g., Butler et al., 2018; Emmons et al., 2012; Dunker et al., 2002; Kwok et al., 2015; Grewe et al., 2013, 2017; Wang et al., 1998; Wu et al., 2011).

Tagging methods are also useful for tracing primary pollutants (e.g., Kranenburg et al., 2013). However, in cases in which the number of different tagged sources is large, tagging methods can become excessively computationally expensive.

In this regard, “adjoint models” (Elbern and Schmidt, 1999; Vautard et al., 2000; Henze et al., 2007) are superior. Adjoint models calculate the derivative of a model scalar with respect to all other model parameters in one single simulation and in this way efficiently quantify the contribution from all emission sources to air pollution in a given recep-

tor region. However, a new adjoint simulation must be performed for each receptor region.

Still, only a relatively small number of sources or receptors at a time can be analyzed by all these methods. Perturbation methods calculate all receptor values for one source group, tagging methods compute all receptors for a limited number of source groups and adjoint models address all sources for one receptor group. Ideally, all contributions to all receptor points should be described.

In this paper we present a method that can efficiently calculate the contribution of a significantly larger number of sources (thousands or more) to a limited (but large) number of receptor regions. This method does not provide results that cannot be obtained by other means, but it does so at a lower computational cost and is thus well suited, especially for time-critical operational applications. It can be built on top of existing Eulerian CTMs relatively easily and thereby has the potential to offer a new range of applications.

An important limitation is that the method is limited to primary pollutants, for which linearity can be assumed. It will thus complement existing methods but not replace them.

In principle the method allows for the definition of any group of sources, but here we will show results only for the case in which each source is defined within a single grid cell. One key limitation, which makes the method manageable, is that the tagged values are stored only up to a preset horizontal and vertical distance from their source. We will call the region within this distance the “local region”. The size of this region must be set as a balance between computational cost and the accuracy requirement of the application.

In the following section we describe the method in technical detail, while in Sect. 3 we show concrete examples of what kind of results the method can provide and how to quantify some of the limitations associated with the method. The results will also be compared against the direct method. In Sect. 4 we will give an overview of what is required to implement the method in an existing CTM and discuss the performance in the EMEP MSC-W implementation. Finally, in the last section we discuss possible applications of the method and plans for further development.

2 Description of the method

In theory the method corresponds to a tagging method, whereby pollutants from different origins are tagged and their values are traced and stored individually. However, the total amount of pollutant is not computed as a sum of tagged values; instead the tagged values show which fraction of the total pollutants originates from a specific origin.

We define the local fraction LF_s in a receptor grid cell as the fraction of pollutant that is due to a particular source term s , and s is the index defining a pollutant or a pollutant from a specific sector. For example, s can refer to primary particulate matter from any sector or be restricted to a power plant

or the road traffic emissions in a specific source region. LF_s is a number between zero and 1 and is calculated as

$$LF_s = \frac{\text{pollutants due to source } s}{\text{total pollutants}} = \frac{LP_s}{TP}. \quad (1)$$

The total pollutant is abbreviated TP; it could be the air concentration of particulate matter, for example. The local pollutant, LP_s , is the part of TP "tagged" from a specific origin s . Its value is in general the result of various processes (emissions, advection, diffusion, etc.) as will be described below. Given the value of the total pollutant, then the local fraction and the local pollutant carry the same information, but as we will see, there are a few practical advantages of storing the local fraction rather than the local pollutant.

In a time-splitting framework the different physical processes are included sequentially, and we will show in the next sections how the value of the local pollutant changes during each of them. For simplicity, the initial value for LF_s is set to zero, given that in the long term LF_s should not be sensitive to the initial value.

2.1 Emissions

The local pollutant and local fraction are associated with a particular emission source category (E_s) in a specific grid cell (formally, the source could also be spread over a group of grid cells, but at present we limit ourselves to sources defined on single grid cells). If $E_s(t)$ is the emission rate of source s at time t , LP_s will increase during the time step Δt :

$$LP_s(t + \Delta t) = LP_s(t) + E_s(t) \Delta t \quad (2)$$

and

$$LF_s(t + \Delta t) = \frac{LP_s(t + \Delta t)}{TP(t + \Delta t)}. \quad (3)$$

For instance, s could refer to emissions of particulate matter from road traffic emissions, TP would be the total concentration of particulate matter in the receptor region, and LF_s would then be the fraction by which the total concentrations in the receptor region would be reduced if the emissions from road traffic in the source region were removed completely (assuming linearity).

2.2 Advection

Transport of pollutants will mix pollutants from different origins. We will individually trace the local pollutant due to different sources and from every horizontal grid cell within the source region. We need then two sets of position indices, one for the origin (source region) and one for the actual position (receptor grid cell):

$$LF_{s,x_s,y_s}(x, y, z, t), \quad (4)$$

where x_s and y_s are the (horizontal) coordinates of the source grid cell, and x, y and z are the coordinates of the receptor

grid cell; s is a specific source category at (x_s, y_s) . In order to keep the calculation at a reasonable cost, one can limit x_s and y_s to be within a preset number of grid cells from the receptor grid cell, Δ^{\max} :

$$x - \Delta^{\max} < x_s < x + \Delta^{\max}, y - \Delta^{\max} < y_s < y + \Delta^{\max}. \quad (5)$$

The source position indices are then replaced by the position relative to the receptor grid cell:

$$LF_{s,\Delta x_s,\Delta y_s}(x, y, z, t), \quad (6)$$

where $\Delta x_s = x_s - x$ and $\Delta y_s = y_s - y$ are the signed distances to the source. In practice, z is also limited, as it is usually not necessary to trace pollutants for receptor grid cells all the way up through the atmosphere. Note that the vertical position of the source is not explicitly traced, but it can, in principle, be included in the form of separate sources s .

We call the region delimited by all (x_s, y_s) and the vertical range of z for the local region.

$LF_{s,\Delta x_s,\Delta y_s}(x, y, z, t)$ is in practice a seven-dimensional array. The range of s depends on the number of source categories to be tracked. The size of this array can be very large, which reflects the large amount of information it carries.

Pollutants can be traced within this region. If they leave the local region, they are no longer identifiable by the method, even if they return to the local region.

Let us consider a flux of pollutant, $F(x, y, z, t)$ (assumed positive), from a grid cell x to $x + 1$ during Δt and a source at a position Δx_s relative to x . The amount of local pollutant leaving the grid cell x is

$$F(x, y, z, t) LF_{s,\Delta x_s,\Delta y_s}(x, y, z, t). \quad (7)$$

At position $x + 1$, the relative position of that source is $x_s - 1$, and the local pollutant is thus updated according to

$$\begin{aligned} LP_{s,\Delta x_s-1,\Delta y_s}(x + 1, y, z, t + \Delta t) = \\ LP_{s,\Delta x_s-1,\Delta y_s}(x + 1, y, z, t) \\ + F(x, y, z, t) LF_{s,\Delta x_s,\Delta y_s}(x, y, z, t). \end{aligned} \quad (8)$$

Or, if the source is moved by one grid cell (Δx_s replaced by $\Delta x_s + 1$), the formula becomes

$$\begin{aligned} LP_{s,\Delta x_s,\Delta y_s}(x + 1, y, z, t + \Delta t) = \\ LP_{s,\Delta x_s,\Delta y_s}(x + 1, y, z, t) \\ + F(x, y, z, t) LF_{s,\Delta x_s+1,\Delta y_s}(x, y, z, t). \end{aligned} \quad (9)$$

The local fractions are then updated according to the definition in Eq. (1).

$$\begin{aligned} LF_{s,\Delta x_s+1,\Delta y_s}(x + 1, y, z, t + \Delta t) = \\ \frac{LP_{s,\Delta x_s+1,\Delta y_s}(x + 1, y, z, t + \Delta t)}{TP(x + 1, y, z, t + \Delta t)} \end{aligned} \quad (10)$$

The fluxes and total pollutants are not explicitly dependent on the source s and are normally available quantities in the CTM.

If the flux is exiting the grid cell x , the local fractions at x do not have to be updated, since it can be assumed that the fractions being removed are the same for the local and total pollutants.

2.3 Diffusion (and convection)

For diffusion we compute the effect of diffusion directly on every local pollutant:

$$LF_{s,\Delta x_s,\Delta y_s}(x,y, :, t + \Delta t) =$$

$$\frac{\text{diffusion}(LP_{s,\Delta x_s,\Delta y_s}(x,y, :, t))}{\text{diffusion}(TP(x,y, :, t))}, \quad (11)$$

where “diffusion()” is the numerical operator that computes the diffusion in the model and the colon indicates its operation over the entire vertical grid column. This ensures a consistent treatment of the local fractions, whatever numerical procedure is applied for the diffusion.

In a practical implementation it is not necessary to include all the vertical levels, as the contribution from higher levels is negligible (it corresponds to pollutants leaving and returning to the local region during the same time step). In our implementation we include only two layers above the highest local region considered.

For convection the same procedure can be used by replacing the diffusion operator in Eq. (11) with the convection operator. In the current EMEP MSC-W model version, the convective processes are not implemented in the local fraction calculations.

2.4 Deposition

For deposition (dry or wet), we can assume that the same fractions of local and total pollutants are removed. Therefore, the local fraction will not vary during the deposition process:

$$LF(t + \Delta t) = LF(t). \quad (12)$$

The simplicity of this formula is one of the motivations for storing LF rather than LP.

2.5 Chemistry

To fully follow the pollutants through all the chemical reactions would, in principle, require an explicit reference to all the sources and grids. It is possible to reduce the size of the problem if linearity is assumed. This has been done by other groups (e.g., Kranenburg et al., 2013). The calculation of all the chemical reactions is one of the most computationally intensive parts of CTMs – roughly 60% in the EMEP MSC-W model (Simpson et al., 2012) used for the tests presented below. A consistent chemical treatment of local pollutants would mean almost multiplying the computation time

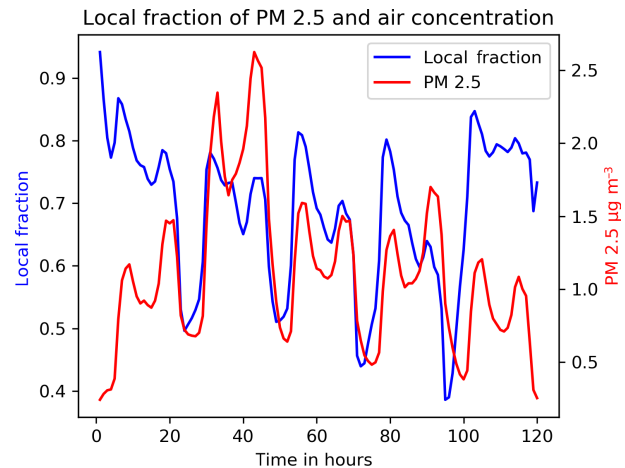


Figure 1. Time evolution of the local fraction of PM_{2.5} in the Oslo agglomeration (left axis) during the period 5 to 9 January 2016 (longitude 11.55°, latitude 59.9°). The total concentration of PM is also shown (right axis).

by the number of local pollutants considered, i.e., the size of $(s, \Delta x_s, \Delta y_s)$. In order to preserve the simplicity of the method, we will in this version assume that the chemical processes modify the local and nonlocal part of the pollutants in the same proportions. With this assumption Eq. (12) can be used. This assumption is correct for primary particles and, as illustrated in our examples below, can also give meaningful results for NH₃, SO_x and NO_x. So far, the method has only been developed for emitted pollutants and not for secondary pollutants.

3 Examples and validation

The local fractions will depend on a broad range of factors such as emission distributions, meteorological conditions, grid resolution, chemical regime and the size of the local region. It is beyond the scope of this article to systematically quantify how all the possible situations affect the local fractions. The limitations of the method should be estimated for each concrete application. The examples in this section also provide methods for estimating different errors associated with the method (limitation of the size of the local region, nonlinearities).

The local fraction $LF_{s,\Delta x_s,\Delta y_s}(x,y,z,t)$ is a seven-dimensional array, and in the following sections we will try to briefly illustrate the information that can be provided by this array.

The results shown in this section are based on a grid with a resolution of 0.3° in the longitude direction and 0.2° in the latitude direction. The parameter settings are essentially the same as what is used for the official EMEP MSC-W model runs, using TNO_MACC-III emissions (2015 update of Kuenen et al., 2014). However, to simplify the interpretation

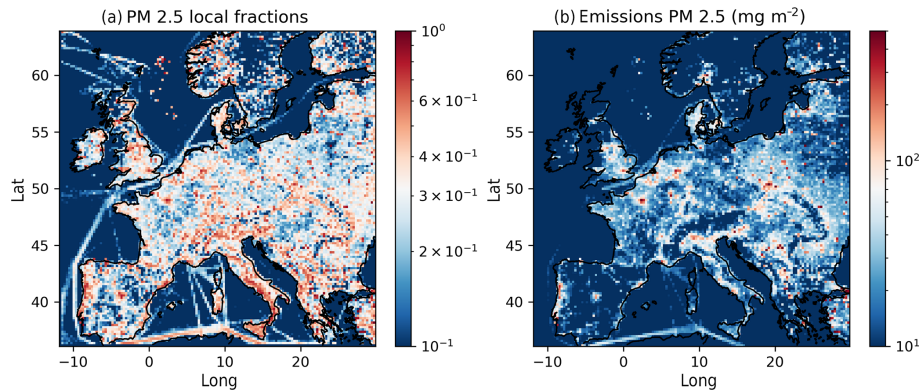


Figure 2. Example of the spatial distribution of the local fraction of $\text{PM}_{2.5}$ averaged over 1 month (March 2016, **a**). The total emissions of $\text{PM}_{2.5}$ accumulated over that period are shown in (**b**).

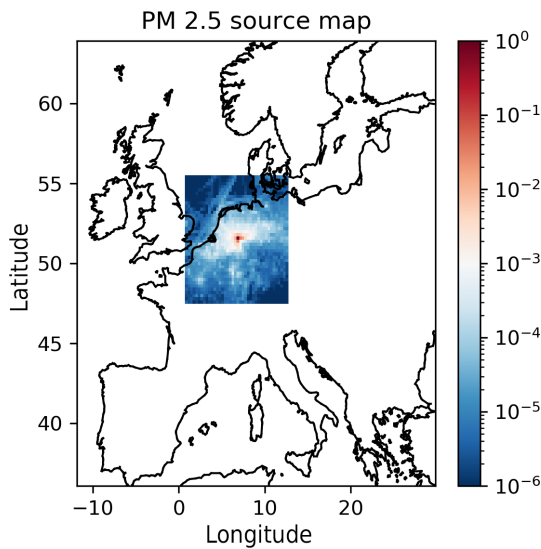


Figure 3. Example of local fractions as a source map. The values show the fraction of $\text{PM}_{2.5}$ that has been emitted at that location and transported to the central point. The sum of all the fractions is in this case 0.976, meaning that 2.4 % of the $\text{PM}_{2.5}$ concentration at the central position originates from sources outside the local region.

of the results, two important modifications have been introduced: a simplified advection scheme is used (see Sect. 3.2), and all emissions are released at the lowest level. The standard settings of the model do not include convection over Europe.

3.1 Time and space dependence ($\Delta x_s = \Delta y_s = 0$)

In Fig. 1, an illustration of the time evolution of the instantaneous local fraction for fine particulate matter ($\text{PM}_{2.5}$) at an arbitrary location (in the Oslo agglomeration) is shown. The value gives the fraction of $\text{PM}_{2.5}$ that has its origin in the same grid cell. It is strongly correlated with the concentrations of $\text{PM}_{2.5}$, but it does not always vary exactly in the

same way. It will also depend on the wind speed, emission rates and the surrounding levels of pollution. If a relatively large amount of clean air is moving into that area, the total concentration will decrease, but the local fraction will remain high. High local fractions indicate that most of the pollutant is locally produced.

Figure 2 shows a map of monthly mean local fractions for March 2016. It gives a picture of how much the sources in a particular grid cell contribute compared to the surrounding sources. The distribution is similar to the emission distribution, but isolated emission sources show up more clearly in the local fractions map.

3.2 Illustration of source–receptor capabilities

For a fixed value of x, y, z and t , the local fractions $\text{LF}_{s,\Delta x_s,\Delta y_s}(x, y, z, t)$ give the contributions of a pollutant s emitted at $(x + \Delta x_s, y + \Delta y_s)$ to the position (x, y) , i.e., a two-dimensional map of the origin of the pollutants found at position (x, y) . Thus, they provide a complete description of all source–receptor relationships within a given distance from the receptor grid cell.

Figure 3 shows such a map for an arbitrary location. It is simply the value of $\text{LF}_{s,\Delta x_s,\Delta y_s}(x, y, z, t)$ averaged over 1 month, where x and y are the position of the central point (receptor). Such a map is calculated for any point on the grid in a single simulation. In this example the local region has a horizontal extent of 41×41 grid cells. Direct methods would then, in principle, require $41 \times 41 + 1 = 1682$ simulations to calculate the values of one of those maps.

In order to compare with the direct method, one can “invert” $\text{LP}_{s,\Delta x_s,\Delta y_s}(x, y, z, t)$ to get a map of the receptors for a fixed source:

$$\text{LP}_{s,\Delta x_s,\Delta y_s}^\dagger(x, y, z, t) = \text{LP}_{s,-\Delta x_s,-\Delta y_s}(x + \Delta x_s, y + \Delta y_s, z, t). \quad (13)$$

$\text{LP}_{s,\Delta x_s,\Delta y_s}^\dagger(x, y, z, t)$ then gives the contributions of a pollutant s located at (x, y) to the position $(x + \Delta x_s, y + \Delta y_s)$.

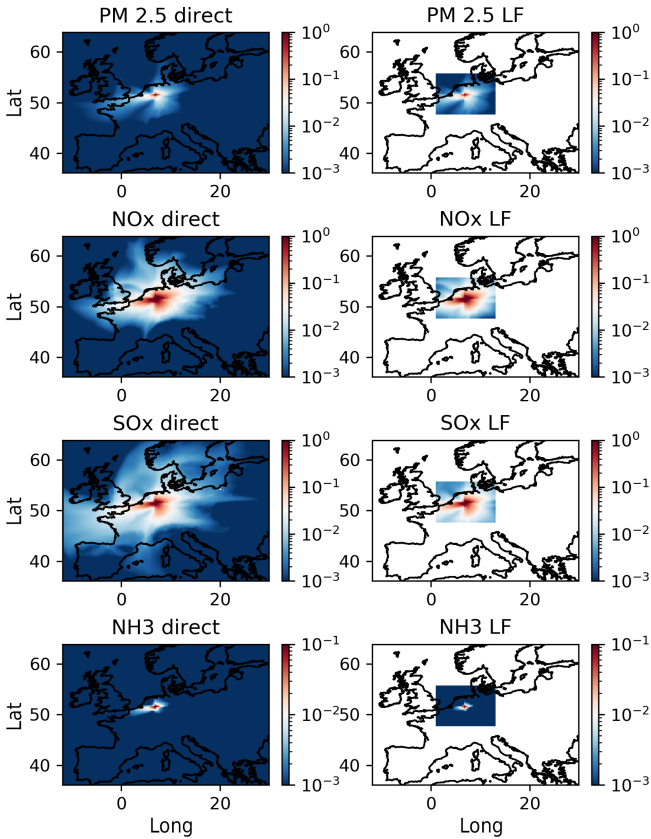


Figure 4. Receptor map for a single grid cell emission obtained through the direct method (left panels) and the local fraction method (right panels) averaged over 1 month (March 2016). Concentrations of PM_{2.5}, NO_x, SO_x and NH₃ (in $\mu\text{g m}^{-3}$). The direct method requires a separate run for each source location. The local fraction method gives the receptor map in one single run in a limited region around the source but for any source grid cell.

Figure 4 illustrates a comparison of the results obtained

1. by removing the emissions from a single grid cell and computing the difference with the normal case (direct method) and
2. by using one single run and Eq. (13) with a local region of size $41 \times 41 \times 8$.

Within the local region the results are similar, but the local fraction method gives such a map for any grid cell in one single run, while the direct method would require a separate run for each source region.

Note that for the purposes of this experiment we have chosen a zero-order advection scheme in all model runs. The default fourth-order scheme is slightly nonlocal, and the direct method would give spurious results very close to the sources; tracking and direct methods would give different results. For example, in the fourth-order scheme, if emissions are *reduced* in one grid cell, this can reduce the flux from the neighboring grid cell in the upwind direction, thereby *in-*

creasing the concentration of pollutants in the upwind grid cell. This is, however, not a problem for the LF method (or any tracking method), and for short distances it is actually an advantage compared to the direct method.

3.3 Vertical transport

For source apportionment applications, the focus is typically on horizontal transport. Nevertheless, the code should trace the pollutants with a combination of vertical and horizontal transport. Over short distances only transport through the lowest layers needs to be considered. If the focus is on regions where most of the pollutants are transported over long distances, the vertical extent of the local area should be chosen large enough.

Figure 5 illustrates the dependence of the local fraction on the thickness of the local region. In this example, only a few vertical levels are required to describe the local fraction within the grid cell (the remaining discrepancy comes from pollutants first leaving the grid cell and then returning later). For a distance of up to 14 grid cells, including eight vertical layers in the local region, results are not distinguishable from the exact value calculated by the direct method. Obviously, emissions or vertical mixing at higher altitudes would require including the corresponding vertical layers.

For NO_x, even for relatively small distances, there is a discrepancy between the contribution calculated with the local fraction method and the direct method (Fig. 6). This is because the local fraction method does not explicitly distinguish between NO and NO₂. The mix modeled in the remote emissions may differ from the local values. Since reaction rates are different for NO and NO₂, the local NO_x transformation rate is not representative for the reaction rates of the incoming “older” NO_x.

3.4 Completeness

For local regions that are large enough, the source of all primary particles can be accounted for. This can be verified directly by summing all the local fractions for a given grid cell:

$$\sum_{\Delta x_s, \Delta y_s} \text{LF}_{s, \Delta x_s, \Delta y_s}(x, y). \quad (14)$$

A sum of 1 means that all sources are accounted for. The difference between the sum of the local fractions and 1 gives the fraction of pollutants with sources outside the local region. In Fig. 7 the sum of the local fractions is shown for every grid cell on the map for different horizontal sizes of the local region. For most land areas, more than 80 % of the sources are found for the smallest window (41×41) and essentially all sources for the largest (161×161).

Figure 8 shows the result for a different vertical extent of the local region. The local fractions are close to complete in most places when eight vertical levels are included (approximately 1522 m height). As one would expect, this roughly

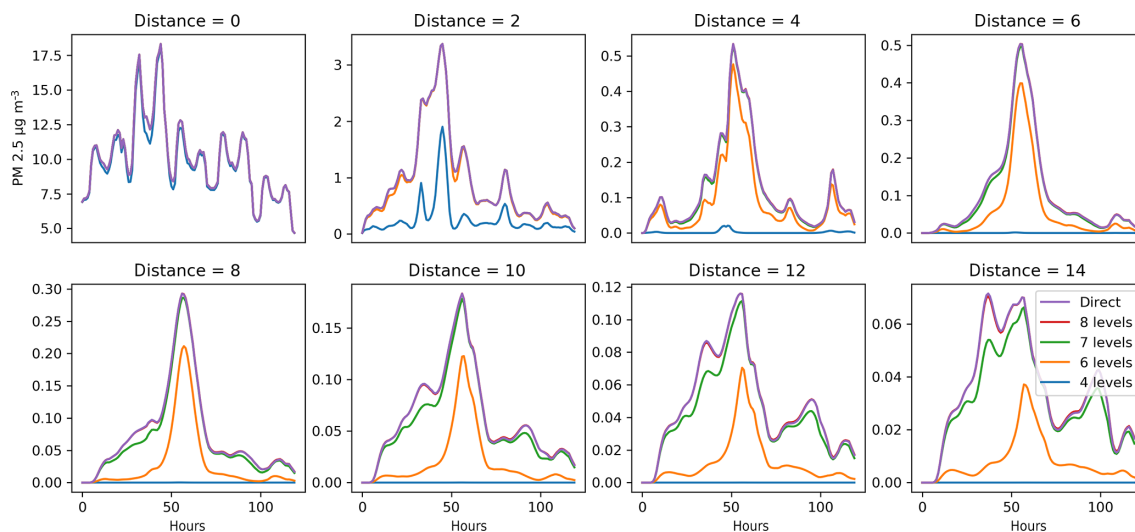


Figure 5. Sensitivity of the concentration of PM_{2.5} (μg m⁻³) to the number of vertical levels included in the local region for different distances from the source. The distance from the source is given in numbers of grid cells (one grid cell represents 0.3° in longitude direction). The source is the Oslo agglomeration. Horizontal axis is time (120 h). The results for eight levels cannot be distinguished from the direct method in the figure, even for the largest distance considered.

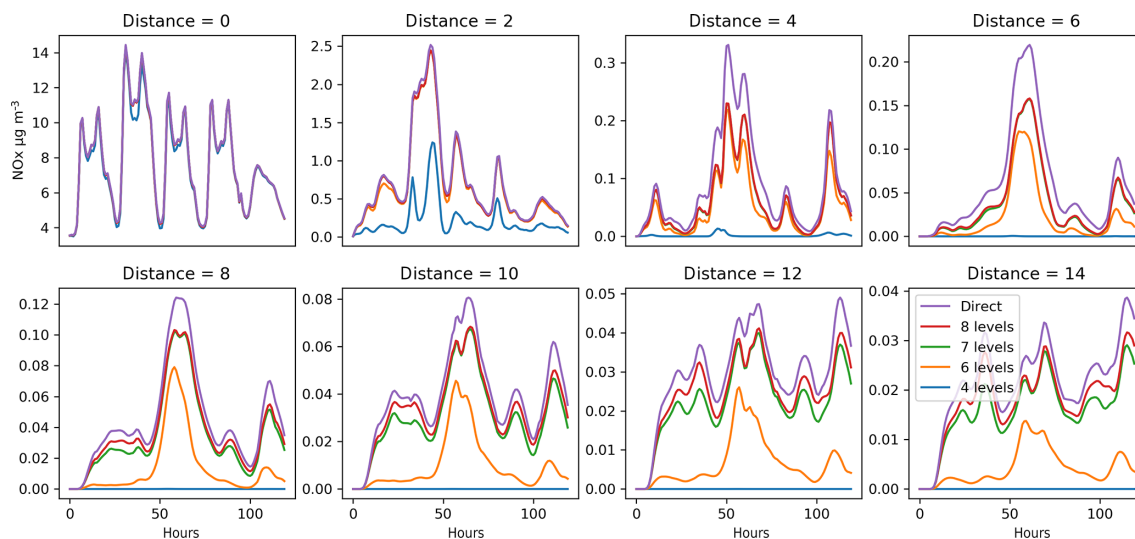


Figure 6. Concentration of NO_x (μg m⁻³). Sensitivity to the number of vertical levels is included in the local region. The distance to the source is given in numbers of grid cells (one grid cell represents 0.3° in longitude direction). The source is the Oslo agglomeration. Horizontal axis is time (120 h). At a distances larger than a few grid cells, a discrepancy can be observed between the contribution calculated with the local fraction method and the direct method.

corresponds to the maximum height of the boundary layer in March over land in Europe.

Note that incomplete results are not a measure of an error in the method. Rather, they show the amount of pollutant with sources outside the local region, which is useful information.

4 Implementation and computational aspects

From an implementation point of view the method is a “diagnostic” calculation in the sense that it gives additional information extracted from existing data, in opposition to a modification of the method for the computation of the concentrations of air pollutants. Therefore, the method can be implemented on top of an existing CTM without having to rewrite the code for the main processes. What is required is including calls to new routines that can perform the operations de-

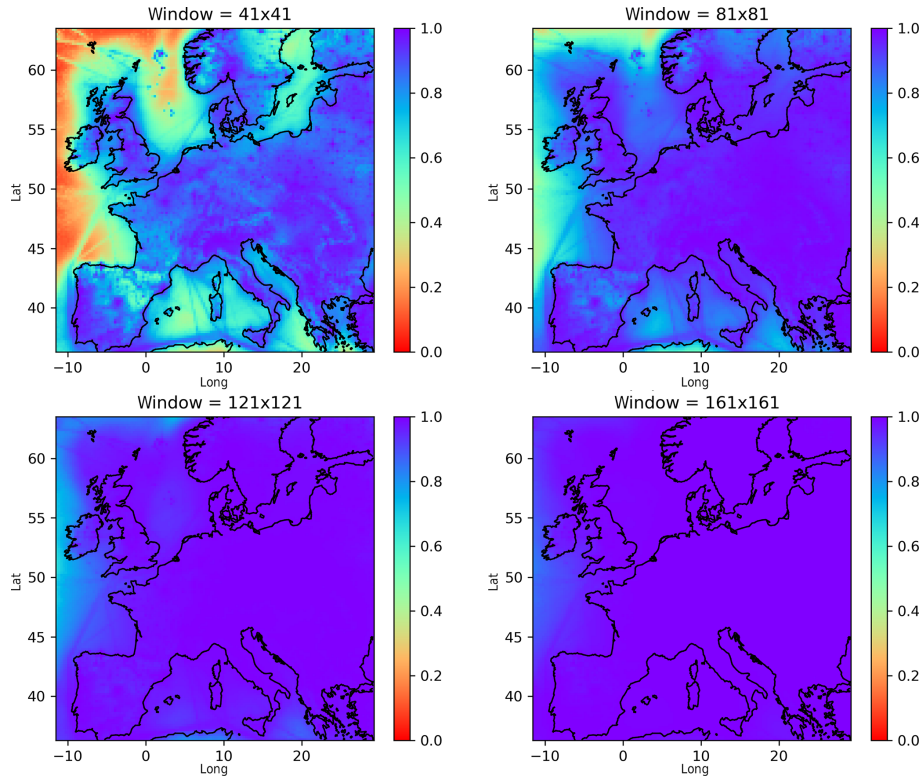


Figure 7. Sum of all local fractions (Eq. 14) for PM_{2.5} and different sizes of the local region (average for March 2016). The distance is counted as the number of grid cells in each direction. All vertical layers (20) are included. A sum of 1.0 means that all the sources have been accounted for.

scribed in Sect. 2. Concretely, the main changes to be made are the following.

- Define the instantaneous local fraction six-dimensional array $LF(s, \Delta x_s, \Delta y_s, x, y, z)$ and one corresponding array for each of the time-averaged periods (at least one for averaging over the run and possibly another for averaging over hours, for example).
- Write a routine that performs the operations from Sect. 2. In addition, there should be a routine for writing out the results (i.e., six-dimensional local fraction arrays), and one routine should do the averaging over time.
- As input for those routines, the main code must make available the emission rates of the relevant sectors and the advection fluxes. If the fluxes are not available or in a simplified version, the fluxes could be defined directly by another method. For example, an already good approximation would be to take $F_x = \frac{c \Delta z}{\Delta x} LP_{up}$, where c is the wind speed in the x direction, Δx the size of the grid cell and LP_{up} the concentrations in the upwind grid cell.
- In Eq. (9) it is necessary to have access to values from the nearest-neighbor grid cells. In parallel implementa-

tions, this may require supplementary communication routines.

- In addition, of course, the calls to those new routines have to be integrated into the main code. Also, switches to choose the pollutants and the sizes of the local region have to be created.

There is no feedback of the LF calculations to the concentrations of air pollutants; those will be unaffected by the new routines. This clear separation greatly simplifies the practical development.

In the EMEP MSC-W implementation (rv4_33), all the extra routines are put in a separate file (uEMEP_mod.f90), except for the LF communication routine. If no LF output is required, those routines are not used at all, and if the LF routines are called, the rest of the code still performs exactly the same operations.

Since one of the key advantages of the local fraction method is its low computational demand, we will give a few concrete examples of the computational cost for providing the local fraction values in our implementation. The transformations carried out for the calculation of local fractions presented in Sect. 2 are all relatively simple. The most computationally intensive parts of the model (calculation of fluxes, chemical transformations, deposition processes) are not ex-

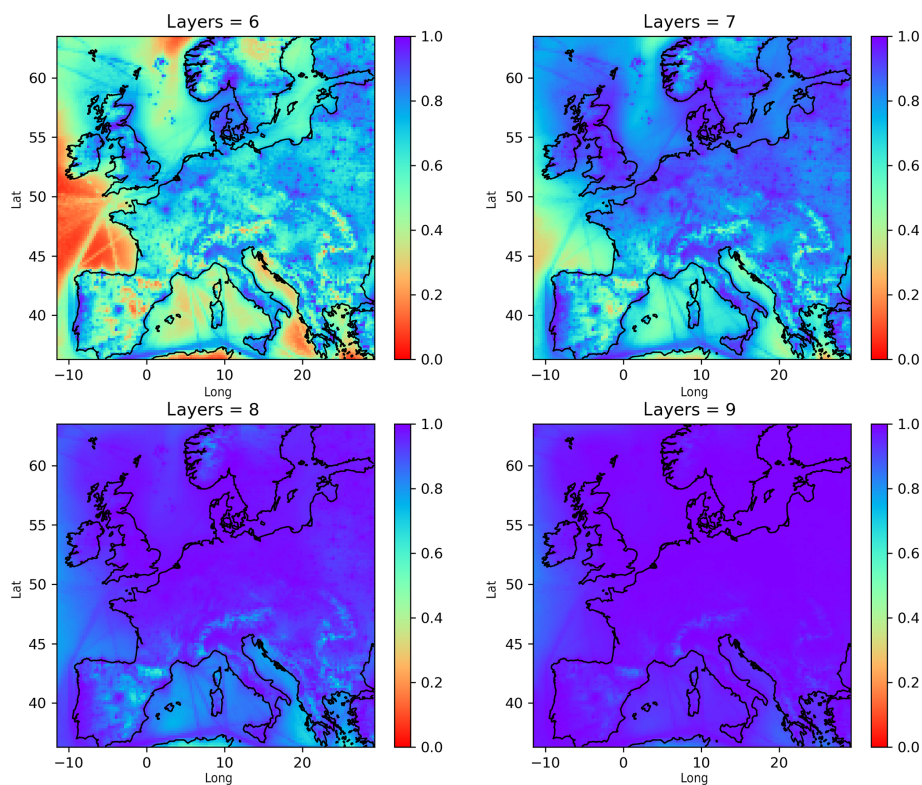


Figure 8. Sum of all local fractions (Eq. 14) for $\text{PM}_{2.5}$ and for different vertical extents of the local region (average for March 2016). A horizontal region of 161×161 is included in the local region. For a standard atmosphere, the heights of the top of layers 6, 7, 8 and 9 are respectively 623, 1015, 1522 and 2149 m. A sum of 1.0 means that all the sources have been accounted for.

Explicitly performed for every local pollutant, but only once for the total concentrations. For processes in which local pollutants are transformed by the same relative amount as nonlocal pollutants (deposition and chemistry in our implementation), there is no need to update the local fractions; this is the main motivation for storing the local fractions rather than the local pollutants.

The calculation of the local fractions only needs information from the nearest neighbors (see Eq. 9) and is therefore well suited for parallel processing in a space partition framework. While storing all the local fractions is memory demanding, the data are distributed among the compute nodes so that the memory requirement can be met by increasing the number of nodes.

In order to illustrate the computational cost, we can consider a typical model run on a $400 \times 260 \times 20$ grid (0.3° longitude \times 0.2° latitude resolution) over 1 month (March 2016) on 160 processors that takes 553 s without the local fraction calculations. Table 1 shows the additional computational cost for computing the local fractions in our implementation. The mathematical operations required to compute the local fractions are proportional to the number of sources considered and the size of the local area. In our implementation the additional time required for advection and diffusion grows faster because of the suboptimal utilization of cache memory. Each

of the processes described in Sect. 2 requires only a few simple mathematical operations on each element of the LF array. The emission part only requires modifying the lowest levels of the array (if emissions are restricted to them). The advection process has to be accounted for in each of the three dimensions and is therefore more costly.

If one is only interested in the nearby sources (within a city, for example), the local fractions can be calculated at almost no additional cost. Remote sources can still be described, but at an additional cost.

A substantial amount of time can be required for writing results to disk if all results are required at finer time resolution, for example every hour (in Table 1 the results are only written out once). This is mainly due to the large amount of data collected; for instance, for a local region of size $81 \times 81 \times 10$ and for each sector or species, $400 \times 81 \times 81 \times 81$ values, which equates to 2.5 GB of data, have to be written to disk each time it is requested (only one vertical level is written out). The corresponding memory demand is calculated in the same way but must further be multiplied by the number of vertical levels of the local region; it must then be multiplied by 2, because one array is needed to store the instantaneous values and one to accumulate the values over time, and multiplied by another factor of 2 because the calculations are done in double precision.

Table 1. Additional computation time needed for the calculation of local fractions in different settings, expressed in percent in comparison to the total time needed when the calculation of local fractions is not included. The first column shows the dimensions of the local region. The second column shows the total additional time required. Columns three to eight show the breakdown of those fractions into the different subroutines (“comm.” stands for communication time between compute nodes). “Other” shows the difference between total time and its components (it is principally due to uncontrolled differences in the speed of the different compute nodes). The last column shows the additional memory required in total; it has to be multiplied by the number of species or sectors requested. The total time without the calculation of local fractions in our test was 553 s.

Local region	Total time	Emission	Advection	Diffusion	Write	Averaging	Comm.	Other	Memory
11 × 11 × 3	1.9 %	0.5	0.8	0.2	0.4	0.1	0.3	−0.4	0.6 GB
21 × 21 × 6	11.9	0.6	5.3	1.3	1.2	0.5	2.4	0.7	4 GB
41 × 41 × 6	38.5	0.9	18.4	4.5	3.8	1.6	9.4	−0.0	16 GB
41 × 41 × 10	63.0	0.9	31.6	6.9	3.8	2.9	17.1	−0.3	26 GB
81 × 81 × 10	241.3	2.0	117.5	27.9	14.3	13.9	67.3	−1.6	102 GB
121 × 121 × 10	623.7	4.7	292.2	108.3	30.2	32.8	170.5	−15.0	227 GB
161 × 161 × 10	1472.1	11.3	740.3	284.4	51.6	62.4	336.1	−14.1	402 GB

5 Discussion

Local fractions are a new concept that can help us understand and analyze the origin of primary pollutants. It has the potential to be developed further, and a new range of applications is still being developed.

Compared to other approaches, there are always trade-offs. The present method cannot at present describe nonlinearities. That excludes all studies of ozone. Long-range transport will also become unpractical at some point, although this is not inherent to the method and could be implemented in the near future.

5.1 Source apportionment

Source–receptor relationships can be produced for any source and receptor within a region around the source. The size of this region can be chosen to be relatively large (100 grid cells or more). Since the fluxes are given from and to individual grid cells, small regions (typically cities) can be studied simply by adding up individual grid cell contributions. These small regions do not have to be predefined in the model simulations. Indeed, the relative contributions of sources that contribute to the pollutants within a city covering several grid cells can be determined in a post-processing step using graphical user interfaces at which the user can choose the source region and source categories interactively.

Still, the method provides information about transport within a limited region only (the local region). The choice of the size of this region is a balance between the computational cost and the distance to the sources of interest. For the study of a city, it may be sufficient to include a region covering the agglomeration. The total pollutants from sources outside the local region are still quantified, but without the specification of their location, using the method presented in Sect. 3.4.

If the goal is to provide source–receptor matrices for large regions (countries), then this method is probably not appropriate in its present form as the computational cost may be

too high, and the level of detail provided is not needed. For such an application the method should be modified so that the tracking is not done for individual grid cells but for larger source areas or group of emission sources.

5.2 Downscaling

One obstacle to combining fine-scale (urban) and regional modeling is the problem of “double counting”. In the regional-scale model, there is usually only one total concentration value, without distinction between its origins. Distinguishing between urban and background pollution can be difficult in practice (Thunis, 2018).

Ideally, the regional model should only compute the background–regional contributions and the fine-scale model can then add the local contribution. In a city, scales down to street level may be required. Those very fine-scale models will not accurately compute the transport between distant streets within the city, and the regional model must account for those. But if the same emission sources are included in both the regional- and fine-scale model, they will be accounted for twice.

The local fractions can give the relative contributions from different sources directly. Thus, it is possible to either redistribute or replace only the appropriate local contributions using the more accurate fine-scale model.

An example for an operational downscaling tool is uEMEP (urban EMEP), which combines the method described in this paper with the EMEP MSC-W air quality model (Simpson et al., 2012) to provide daily air quality forecasts for all of Norway (<https://luftkvalitet.miljostatus.no/>, last access: 25 March 2020; Denby et al., 2020).

5.3 Improved modeling

Concentrations of pollutants near the surface are required to assess health impacts or dry deposition. However, in many CTMs, the lowest layer is several tens of meters thick, and

the concentrations of pollutants will have a nonconstant vertical profile within the layer. The shape of the profile will depend on the local conditions: if the pollutants are emitted locally at the surface the concentration will typically decrease with height, while the opposite is true for background pollutants. With knowledge of the local fractions it is possible to improve the description of the vertical profile and thus obtain a more accurate estimation of, for instance, 3 m concentrations (useful for health impact studies) or dry deposition rates.

As shown in Figs. 1 and 2 the local fractions vary strongly in space and time. If this information can be used to give better estimations of the vertical profiles of pollutants it should have a significant effect on the results.

5.4 Future work

In this work, sources are always defined in an individual grid cell. The relative position of the source, $(\Delta x_s, \Delta y_s)$, could be replaced by a generic index that would point to more general groups of grid cells or regions. The formalism would be the same, except that emissions from any grid cell from the relevant region should be added together in the local fraction. This would allow, for instance, for individually distinguishing all grid cells in the immediate vicinity of the receptor grid cell and successively larger regions as the distance increases. Another application could be to define countries as emitter regions.

In the future we plan to generalize the method to also include chemical processes in some simplified form. The ambition is to still provide information for a very large number of sources but to describe chemical processes in an approximate way. Compared to existing tagging methods, it will trade accuracy for computational efficiency.

Code and data availability. The full EMEP MSC-W model code and main input data are publicly available through a GitHub repository under a GNU General Public License v3.0 (name emep-ctm). The routines related to the local fractions are part of the standard model. The exact version of the model used to produce the illustrative examples used in this paper (rv4.33) is archived on Zenodo (<https://doi.org/10.5281/zenodo.3265912>, Valdebenito et al., 2020).

Author contributions. All authors contributed to the discussion and development of the main ideas, the applications, and the preparation of the paper. PW wrote the corresponding Fortran code.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The computations were performed on resources provided by UNINETT Sigma2 – the National Infrastructure for High Performance Computing and Data Storage in Norway.

Financial support. This research has been supported by the AirQuip project funded by Norges Forskningsråd (grant no. 267734).

Review statement. This paper was edited by Patrick Jöckel and reviewed by two anonymous referees.

References

- Butler, T., Lupascu, A., Coates, J., and Zhu, S.: TOAST 1.0: Tropospheric Ozone Attribution of Sources with Tagging for CESM 1.2.2, *Geosci. Model Dev.*, 11, 2825–2840, <https://doi.org/10.5194/gmd-11-2825-2018>, 2018.
- Clappier, A., Belis, C. A., Pernigotti, D., and Thunis, P.: Source apportionment and sensitivity analysis: two methodologies with two different purposes, *Geosci. Model Dev.*, 10, 4245–4256, <https://doi.org/10.5194/gmd-10-4245-2017>, 2017.
- Denby, B. R., et al.: *Geosci. Model Dev.*, in preparation, 2020.
- Dunker, A. M., Yarwood, G., Ortmann, J. P., and Wilson, G. M.: Comparison of Source Apportionment and Source Sensitivity of Ozone in a Three-Dimensional Air Quality Model, *Environ. Sci. Tech.*, 36, 2953–2964, <https://doi.org/10.1021/es011418f>, 2002.
- Elbern, H. and Schmidt, H.: A four-dimensional variational chemistry data assimilations scheme for Eulerian chemistry transport modeling, *J. Geophys. Res.*, 104, 18583–18598, 1999.
- Emmons, L. K., Hess, P. G., Lamarque, J.-F., and Pfister, G. G.: Tagged ozone mechanism for MOZART-4, CAM-chem and other chemical transport models, *Geosci. Model Dev.*, 5, 1531–1542, <https://doi.org/10.5194/gmd-5-1531-2012>, 2012.
- Fiore, A. M., Dentener, F. J., Wild, O., Cuvelier, C., Schultz, M. G., Hess, P., Textor, C., Schulz, M., Doherty, R. M., Horowitz, L. W., MacKenzie, I. A., Sanderson, M. G., Shindell, D. T., Stevenson, D. S., Szopa, S., Van Dingenen, R., Zeng, G., Atherton, C., Bergmann, D., Bey, I., Carmichael, G., Collins, W. J., Duncan, B. N., Faluvegi, G., Folberth, G., Gauss, M., Gong, S., Hauglustaine, D., Holloway, T., Isaksen, I. S. A., Jacob, D. J., Jonson, J. E., Kaminski, J. W., Keating, T. J., Lupu, A., Marmor, E., Montanaro, V., Park, R. J., Pitari, G., Pringle, K. J., Pyle, J. A., Schroeder, S., Vivanco, M. G., Wind, P., Wojcik, G., Wu, S., and Zuber, A.: Multimodel estimates of intercontinental source-receptor relationships for ozone pollution, *J. Geophys. Res.*, 114, D04301, <https://doi.org/10.1029/2008JD010816>, 2009.
- Folberth, G. A., Rumbold, T. S., Collins, W. J., and Butler, T. M.: Global radiative forcing and megacities, *Urban Climate*, 1, 4–19, 2012.
- Grewe, V.: A generalized tagging method, *Geosci. Model Dev.*, 6, 247–253, <https://doi.org/10.5194/gmd-6-247-2013>, 2013.
- Grewe, V., Tsati, E., Mertens, M., Frömming, C., and Jöckel, P.: Contribution of emissions to concentrations: the TAGGING 1.0 submodel based on the Modular Earth Submodel System (MESSy 2.52), *Geosci. Model Dev.*, 10, 2615–2633, <https://doi.org/10.5194/gmd-10-2615-2017>, 2017.

- Henze, D. K., Hakami, A., and Seinfeld, J. H.: Development of the adjoint of GEOS-Chem, *Atmos. Chem. Phys.*, 7, 2413–2433, <https://doi.org/10.5194/acp-7-2413-2007>, 2007.
- Jacob, D. J., Logan, J. A., and Murti, P. P.: Effect of rising Asian emissions on surface ozone in the United States, *Geophys. Res. Lett.*, 26, 2175–2178, 1999.
- Kranenburg, R., Segers, A. J., Hendriks, C., and Schaap, M.: Source apportionment using LOTOS-EUROS: module description and evaluation, *Geosci. Model Dev.*, 6, 721–733, <https://doi.org/10.5194/gmd-6-721-2013>, 2013.
- Kwok, R. H. F., Baker, K. R., Napelenok, S. L., and Tonnesen, G. S.: Photochemical grid model implementation and application of VOC, NO_x, and O₃ source apportionment, *Geosci. Model Dev.*, 8, 99–114, <https://doi.org/10.5194/gmd-8-99-2015>, 2015.
- Kuenen, J. J. P., Visschedijk, A. J. H., Jozwicka, M., and Denier van der Gon, H. A. C.: TNO-MACC-II emission inventory; a multi-year (2003–2009) consistent high-resolution European emission inventory for air quality modelling, *Atmos. Chem. Phys.*, 14, 10963–10976, <https://doi.org/10.5194/acp-14-10963-2014>, 2014.
- Simpson, D., Benedictow, A., Berge, H., Bergström, R., Emberson, L. D., Fagerli, H., Flechard, C. R., Hayman, G. D., Gauss, M., Jonson, J. E., Jenkin, M. E., Nyíri, A., Richter, C., Semeena, V. S., Tsyro, S., Tuovinen, J.-P., Valdebenito, Á., and Wind, P.: The EMEP MSC-W chemical transport model – technical description, *Atmos. Chem. Phys.*, 12, 7825–7865, <https://doi.org/10.5194/acp-12-7825-2012>, 2012.
- Thunis, P.: On the validity of the incremental approach to estimate the impact of cities on air quality, *Atmos. Environ.* 173, 210–222, 2018.
- Thunis, P., Clappier, A., Tarrason, L., Monteiro, A., Pisoni, E., Wesseling, J., Belis, C. A., Pirovano, G., Janssen, S., Guerreiro, C., and Peduzzi, E.: Source apportionment to support air quality planning: Strengths and weaknesses of existing approaches, *Environ. Int.*, 130, 104825, <https://doi.org/10.1016/j.envint.2019.05.019>, 2019.
- Valdebenito, A., Wind, P., and Benedictow, A.: metno/emep-ctm: OpenSource rv4.33 (201906), <https://doi.org/10.5281/zenodo.3265912>, last access: 25 March 2020.
- Vautard, R., Beekmann, M., and Menut, L.: Applications of adjoint modelling in atmospheric chemistry: sensitivity and inverse modelling, *Environ. Modell. Softw.*, 15, 703–709, 2000.
- Wang, Y., Jacob, D. J., and Logan, J. A.: Global simulation of tropospheric O₃-NO_x-hydrocarbon chemistry: 3. Origin of tropospheric ozone and effects of nonmethane Hydrocarbons, *J. Geophys. Res.*, 103, 10757–10767, 1998.
- Wu, Q. Z., Wang, Z. F., Gbaguidi, A., Gao, C., Li, L. N., and Wang, W.: A numerical study of contributions to air pollution in Beijing during CAREBeijing-2006, *Atmos. Chem. Phys.*, 11, 5997–6011, <https://doi.org/10.5194/acp-11-5997-2011>, 2011.