The Lagrangian particle dispersion model FLEXPART version 10.3: Response to reviewers

Ignacio Pisso[1], Espen Sollum[1], Henrik Grythe[1], Nina I. Kristiansen[1,*], Massimo Cassiani[1], Sabine Eckhardt[1], Delia Arnold[2,3], Don Morton[4], Rona L. Thompson[1], Cristine D. Groot Zwaaftink[1], Nikolaos Evangeliou[1], Harald Sodemann[5], Leopold Haimberger[6], Stephan Henne[7], Dominik Brunner[7], John F. Burkhart[8], Anne Fouilloux[8], Jerome Brioude[9], Anne Philipp[6,10], Petra Seibert[11], and Andreas Stohl[1]

[1]Norwegian Institute of Air Research, Kjeller, Norway
[2]Central Institute for Meteorology and Geodynamics, Vienna, Austria.
[3]Arnold Scientific Consulting, Manresa, Spain
[4]Boreal Scientific Computing, Fairbanks, Alaska, USA
[5]Geophysical Institute, University of Bergen, Norway
[6]Department of Meteorology and Geophysics, University of Vienna, Austria
[7]EMPA, Dübendorf, Switzerland
[8]Department of Geosciences, University of Oslo, Norway
[9]LACy, Université de la Réunion, France
[10]Aerosol Physics & Environmental Physics, University of Vienna, Austria
[11]University of Natural Resources and Life Sciences, Institute of Meteorology, Vienna, Austria
[*]now at Met Office, FitzRoy Road, Exeter, EX1 3PB, UK

1 Anonymous Referee #1

Received and published 26 March 2019.

FLEXPART is a well-known Lagrangian particle dispersion model that has been used in various applications. The new version presented here has some very important up- dates such as considering the skewness in the vertical velocity distribution and adding option to start backward runs from depositions. The manuscript clearly described the new features and provided detailed information for the usage of the whole software package. It is well organized and presented.

1.1 Major points:

1.1.1

1) In parallelisation section, the test examples have 40 million particles released. It is not quite practical to have that many particles in some applications. As it has been pointed out, the speedup highly depends on the number of particles released. Can the authors add examples with less particles (less than a million)?

We provide a practical small-scale example (that can be run on a laptop) in the new section 3.2.3 and new Table 6 with 900 000 particles. Such test case takes about 1 minute per simulation day on 4 cores. (for 15 days it takes about 1hour to run on 1 core and 20 minutes on 4 cores.)

1.1.2

2) In section 3.3, the authors commented that a direct comparison between serial and parallel version is difficult. Is it a viable way to specify random seed numbers explicitly in the two versions?

The random seeds are hard coded but can be changed. However, changing the random seeds would not make direct comparison serial/parallel easier when running running in multiple cores. In order to remove any statistical uncertainty, each numerical operation would need to be executed with exactly the same random number value,

which is impractical in an MPI setting. However, as long as the difference in the model output between serial and parallel versions is within the statistical error there is no reason to assume that results are essentially different. It is written (2nd paragraph) that "To confirm that the only source of differences between the serial and parallel code is in the random number generation, we first observe that when the parallel executable is run using only one process, it produces results identical to the serial version." ... "The outputs from the serial and parallel versions of the code when run this way are identical except for small differences due to round-off errors". The difference is statistical in parallel calculations with different number of cores because each core uses its own random stream.

We have clarified the paragraph in order to express more clearly that the difference is only statistical: "In order to ensure that the parallel version produces results with the same accuracy as the serial version, we have performed a set of tests and validation experiments. A direct comparison between the versions can only be performed in statistical terms because FLEXPART uses Gaussian distributed random numbers for calculating turbulent velocities of the particles. For the parallel version we let each process independently calculate a set of random numbers, which leads to small numeric differences (arising from the random 'noise') between the versions."

1.1.3

3) Figure 9 might be redundant as the information has been well described in text.

We have removed Figure 9 and left the explanation in the text.

1.2 Minor points:

1.2.1

Page 1, lines 5-6: Note that "volcanic emissions" are not "atmospheric gases and aerosols" for which the examples are given. Consider to replace it with "volcanic ash" or something else.

It is now replaced with "volcanic ash".

1.2.2

Page 2, line 3: "since many years" -> "for many years" or "since many years ago"

The wording of the phrase was changed.

1.2.3

Page 5, line 33: It is not accurate to state that all Eulerian models have the tracer instantaneously mixed within a grid box although most of them do.

We clarified that most of them do.

1.2.4

Page 6, line 3: The term "air history" is confusing. It should be changed.

We changed "air history" for "can also be used to investigate the history of air parcels affecting designated atmospheric point measurements".

1.2.5

Page 7, line 22: Fluxes across grid cell center "lines"? Are they supposed to be "faces"?

We replaced "grid cell center lines" with "grid cell faces".

1.2.6

Page 11, Figure 1: The figure is not easy to read. Can it be changed to color-filled contour flood style?

The figure was changed to color-filled contour flood style.

1.2.7

Page 15, Equation (7): It is easy to mistake iC_r as two variables, i and C_r. Can it be changed?

We defined it as ic_r in order to keep the notation consistent with Grythe et al. (2017) where it was originally defined.

1.2.8

Page 21, line 1, ”For a, say, 10-day ... ”: It is better to replace it with ”For instance, a 10-day ...”.

The phrase was changed (“For instance, for a 10-day backward simulation... ”).

1.2.9

Page 15, line 15: It is not clear how the emissions are defined here. Are they assumed constant over the entire simulation period?

We couldn’t associate the page and line with a particular simulation. Unless specified otherwise, the emissions are assumed constant over the entire simulation period. Simulation shown in Fig 2 uses variable source term described in Stohl et al. (2012).

1.2.10

Page 24, Figure 5: Is the unit of molecules $cm^{-3} \times 10^6$ considered mass based? It would be good to add a sentence or two to comment on units like this.

The unit molecules cm^{-3} is related to the mass concentration simply by dividing the number of molecules by Avogadro’s number to get the number of moles of OH, and multiplying the number of moles by the molecular weight of OH (i.e., 17). This gives the unit $g\,cm^{-3}$.

1.2.11

Page 30, Table 4: It is strange to have the parallel efficiency greater than 1. What is the reason for that? Are multiple realizations needed to have robust numbers? In addition, It might be better to leave the column with 1 process blank than to fill it with 1.00 or 1.000.

The parallel efficiency may be greater than 1 because of superlinear speedup (Ristov et al., 2016). Superlinearity can be a consequence of memory usage (cache effect, changes in pagination) and storage. Running on 2 cores instead of 1 effectively doubles the amount of “fast” cache memory available. Another possible source are changes in network latency. We have clarified the factor and added a reference in the text.

1.2.12

Page 38, line 9: Is there a typo here with “LOWER release height ABOVE 0.5 m”?

The typo was corrected.

1.2.13

Page 43, Table 9: It is better to have ”mixing ratio” appear in the description for ”grid_pptv_date_nnn” (row 6).

The description was modified.

1.2.14

Page 44, line 19: The file name has been cut.

The overfull in LaTeX was corrected by rearranging the phrase.

1.2.15

Page 45, Table 10: The values of 3 and 4 are shown here for `ind_receptor`, but they are not listed in Table 6 as possible values. They need to be consistent.

The values of 3 and 4 were added to line 23 of table 6 in order to be consistent.

1.2.16

Page 47, line 6: Flexpart -> FLEXPART Page 47, line 30: Group 10 -> Group ten

The reference to the website was corrected. The number 10 was replaced with “ten”.

---

2 Anonymous Referee #2

Received and published: 14 May 2019

The paper describes FLEXPART model development between versions 6.2 (last documented in Stohl et al. (2005)) and 10.3. It is well written, well presented and the authors have clearly spent significant time ensuring that it is ready for publication. A number of minor comments, suggestions and corrections are listed below. I have just one main comment concerning what is and what is not model development and where the line falls between FLEXPART and pre- and post-processing code and visualisation tools. The paper contains details of a number of (clearly very practical) uses and external code (e.g., the use of FLEXPART to calculate SRR, code to generate emissions and input meteorological files). It is beneficial to the community to share these for wider use. My question is, what constitutes FLEXPART and what is not part of the model? And, is the paper a FLEXPART model development paper or broader than this?

We thank the reviewer for the remark, so we can provide a clarification. FLEXPART consists of the (Fortran) source files required to build an executable, not including external libraries such as eccodes etc. The makefiles and the sample input as provided in the "options" may also be included under this term. However, in order to do real work with FLEXPART, one needs also to obtain meteorological input data (and in the case of ECMWF, this is not trivial and the `flex_extract` package is provided for this), and one needs to do postprocessing. This is the reason why we include a selection of such tools here.

2.1 Minor Issues:

2.1.1

1. Simulations representing material everywhere (such as background fields), referred to in this paper as 'filling the whole atmosphere' or 'domain-filling mode', are particularly expensive in Lagrangian particle models. There must be some limit on the scale of such simulations which can be feasibly be conducted. The authors give no indication of this issue and furthermore, 'filling the whole atmosphere' (page 6, line 17) or 'the entire atmosphere is represented by particles' (page 6, line 33) gives the impression of global simulations. Can the authors discuss this? An example, giving the number of particles used and the size of the computational domain for such a simulation, would also assist here.

Simulations for the whole atmosphere can be carried out with a number of particles that allows FLEXPART to run in a standard linux workstation. Stohl (2006) employed 1.4 million particles during 5.5 years for global simulations while investigating Arctic transport. Pisso et al. (2010) employed 2.2 million trajectories in global calculations simulating stratospheric residence times in order to investigate of tropospheric transport into the stratosphere. Simulations for 20 years were run in a single workstation. Disk space will limit the availability of input wind fields as well as the output frequency. The simulation time clearly depends on the simulated period and multi year simulations can take several days. In a cluster, where large databases of winds are available, it is routine to run multi decadal simulations (e.g. stratospheric ozone or water cycle experiments). Although possibly at a lower resolution and taking into account fewer process than a full Eulerian CTM/GCM, such simulations may provide valuable insights on the pathways and history of air masses in Earth's atmosphere. A clarification was added to the text: "FLEXPART can also be used in a domain-filling mode where the entire atmosphere is represented by (e.g. a few millions of) particles of equal mass (Stohl and James, 2004). The number of particles required for domain-filling simulations, not unlike those needed for other types of simulations, depends on the scientific question to be answered. For instance, a few million particles distributed globally are often enough to investigate statistical properties of air mass transport (e.g., monthly average residence times in a particular not too small area) but would not be enough for a case study of airstreams related to a particular synoptic situation (e.g., describing flow in the warm conveyor belt of a particular cyclone)."

2.1.2

2. Information on the resolution of the meteorological data commonly used (e.g., the current temporal and spatial resolution of the ECMWF IFS and NCEP GFS input data) would be of interest and give some context (page 7, initial paragraph).

Common spatial resolutions for IFS depending on the application include $0.75^\circ \times 0.75^\circ$, 3 h (standard for older products, e.g. ERA Interim), $0.4^\circ \times 0.4^\circ$, 1 h (standard for newer products, e.g. ERA5), and $0.125^\circ \times 0.125^\circ$, 1 h (current ECMWF operational data). The model has currently 137 vertical levels. NCEP GFS input files are usually used at $1^\circ \times 1^\circ$ horizontal resolution, 64 vertical levels and 6 h time resolution. NCEP GFS input files are also available at $0.5^\circ \times 0.5^\circ$ and $0.25^\circ \times 0.25^\circ$ horizontal resolution. This additional information was added to the text.

2.1.3

3. What is meant by 'input data' (page 7, line 15)? Is it just meteorological input data?
In this case "meteorological input data" is meant', this was clarified.

2.1.4

4. Typo in Figure 1 caption: The latitude and longitude values have been interchanged. The source is at 16.3274 degrees N (not E) and vice versa.
We checked that the source is near Vienna (Austria) at 47.9157 N and 16.3274 E. The figure is integrated in the latitude axis (i.e. longitude-height).

2.1.5

5. The authors mention that changes in the settling velocities are largest in the cold upper troposphere and this will impact the residence time of volcanic ash particles. Out of curiosity, does the new scheme increase or decrease the residence time?
The residence times in the upper troposphere are increased with the new scheme, but the effect is not particularly large, typically on the order of 20%.

2.1.6

6. On page 12 (line 30) it is stated that wet scavenging does not occur when the precipitation rate does not exceed 0.01 mm/hr. Some NWP models have had a problem with excessive amounts of drizzle. Is this limit related to this?
This limit is related to the empirical fits of the observed scavenging coefficients to precipitation rate reported by Laakso et al. (2003) and Kyroe et al. (2009). These fits are not valid for extremely low precipitation rates (some scavenging would occur even in the absence of precipitation) and we thus cannot use these parameterization for the very lowest precipitation rates. The impact, however, should be small, as little scavenging occurs under these conditions. The limit is indeed intended to also catch spurious numerical precipitation. Related to this comment, we found that our description was not 100% accurate. While we wrote that scavenging is not applied when the total precipitation should not exceed 0.01 mm/h, it is actually not applied if neither the large scale nor the convective precipitation rate exceeds 0.01 mm h^{-1}. This has been corrected in the manuscript.

2.1.7

7. It might help the reader to add '(liquid and/or ice)' to '3-D cloud water mixing ratio' (lines 22-23, page 13) to clarify that it includes both components.
NCEP only provides the total of liquid and ice cloud water. To clarify this, we have added in brackets: (liquid plus ice).

2.1.8

8. It's not clear what 'it' refers to (line 8, page 14). 'Sub-grid area fraction' or 'sub-grid precipitation rate' or 'grid-scale precipitation rate'?
The entire parameterization is meant. We have changed "It" to read "This parameterization of sub-grid variability".

2.1.9

9. There is an inconsistency in the definition of alpha in equation 4. The text says that for $T > T_L$, $alpha = 1$, whereas in equation 4, $alpha = 0$ for $T = T_L$. Likewise for $T < T_I$, the text has $alpha = 0$ whereas equation 4 indicates it should be $alpha = 1$. For equation 4 to agree with the text, I think the T_L on the numerator should be T_I. Equation 9 also implies alpha is the ice fraction rather than the liquid fraction. I suspect, that equation 9 should be $F_{nuc} = alpha * CCN_{eff} + (1 - alpha) * IN_{eff}$.
We thank the reviewer for pointing this out. We corrected in the text the definition of F_{nuc} (the fraction of aerosols that reside within cloud droplets) and the misplaced T_L instead of T_I. Unfortunately, the error in the definition of alpha was also present in the model code. So we are even more thankful to the reviewer for catching this issue, which escaped the attention of all co-authors. For clarity, we introduced the variables α_L and α_I for liquid and ice fractions of the cloud water respectively. The definition of α_I is consistent with that in Figure 2 of

Grythe et. al (2017). In order to assess the impact of the change on the results we performed a series of additional tests with monthly FLEXPART simulation for the Zeppelin station based on the CERA wind fields. Four backward simulations were performed. Two with the original version, two with the corrected version. For each version, we calculated in one simulation the deposited mass and in the other simulation the surface concentration. The BC concentrations were estimated from the annual average of BC emissions. Figure 1 shows that the impact of the change is of the order of 15% in this application case. In application cases where the values of CCN_{eff} and IN_{eff} are assumed to be identical, the error has no effect. This was, for instance, the case for the calculations shown in Figure 2 of the manuscript.

2.1.10

10. Is the cloud and precipitation data time averaged or instantaneous (lines 17-24, page 14)? Peculiarities in wet deposition fields are obtained when using instantaneous data at high spatial resolution but with insufficient temporal resolution (of the input meteorological fields).

ECMWF precipitation data as available from the ECMWF archives are accumulated over the forecast time period but, after extraction from the archives, are de-accumulated by the flex_extract software so that their validation times are the same as for the other, instantaneous meteorological data (e.g., winds). Despite de-accumulation, the precipitation data must be considered temporal averages. Cloud data (liquid and ice water content, or their sum), on the other hand, are instantaneous quantities directly extracted from the ECMWF archives without de-accumulation. The reviewer is right that this causes inconsistencies between the two meteorological data sets used in the wet scavenging parameterizations, and the use of instantaneous data may cause specific problems if the temporal resolution is insufficient (the same could also be said for the wind data, though). This has been discussed in detail by Hittmeir et al. (2018) and we are planning to implement their improved interpolation scheme in the near future. However, even that will not entirely resolve this issue. In principle, the use of time-averaged quantitities for all meteorological data would be best, but such data are currently not available from ECMWF.

2.1.11

11. Is cl (lines 17, 21, equation 8, page 15) the same as c_l (line 27, equation 2, page13)? If so, I suggest consistent notation. If not, then cl needs to be defined here. I would expect S_i to increase with increasing cloud water and this seems contrary to equation 8, if cl is cloud water.

Thank you for discovering this inconsistency. Yes, cl and c_l are the same. This has been corrected. The scavenging is assumed to be proportional to the cloud water removal rate I_s/c_l (see also Equation 7 in the paper). Ideally, of course, one would want to specify the removal rate as a function of altitude. However, since precipitation (I_s) is only a two-dimensional variable in the available ECMWF data, this is not easily possible and would require additional assumptions. We therefore applied a vertically homogeneous scavenging.

2.1.12

12. CCN and IN efficiencies 'increase with increasing particle size', is this true? Good cloud condensing nuclei are sub-micron in size.

It is true that CCNs are typically of sub-micron size, simply because these particles are the only ones frequent enough to be of importance. However, based on Köhler theory, a certain particle size is needed in order to overcome the energy barrier of homogeneous nucleation. Generally, the activated fraction of particles of the same type increases with particle size - this is most relevant for particles smaller than about 200 nm. See, for instance, Anttila et al. (2012) for a detailed investigation of the size dependency.

2.1.13

13. Where is Appendix 5.1.3 (line 24, page 16)?

Section 5.1.3 (description of SPECIES files) was meant. This was corrected.

2.1.14

14. Line 13 (page 18) doesn't make sense ?. . .at the example. . .??

We rewrote the phrase.

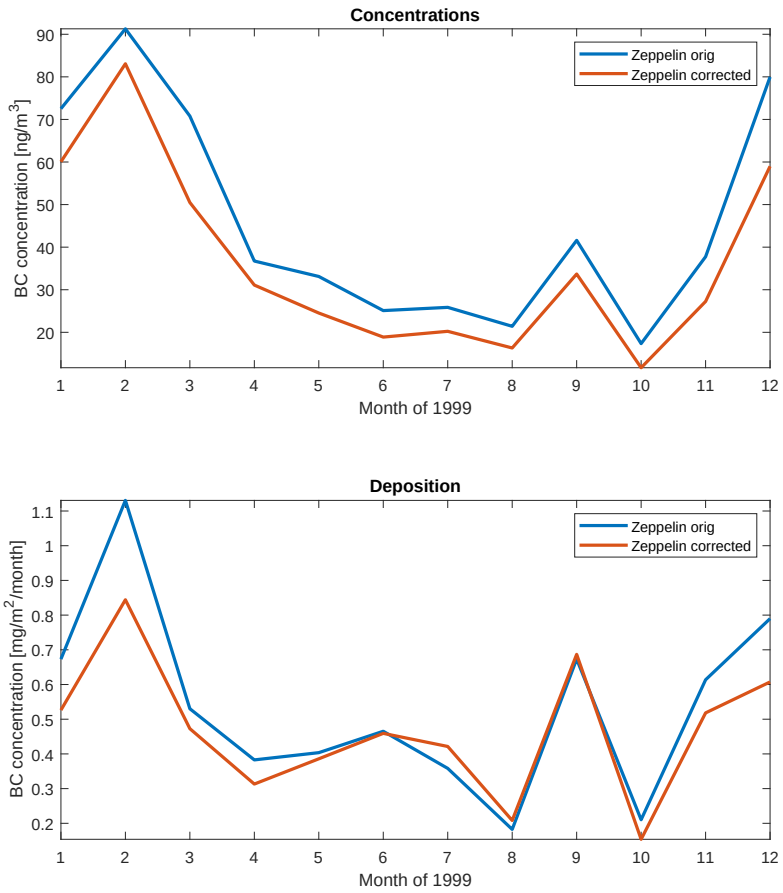

Figure 1: Deposited mass by wet deposition and surface concentration in the lowest level before and after correcting the error in equation 4

2.1.15

15. Can the observation location be marked in Figure 3?
We have modified the plot. The marker is included in the final version of the manuscript.

2.1.16

16. The sentence which spans pages 20 and 21 doesn't seem to follow across the page break. Should Seibert and Frank be in parentheses?
The parentheses were added.

2.1.17

17. MPI has recently been introduced into NAME (version 8) (line 1, page 27).
We changed the description.

2.1.18

18. It would help to indicate that Table 3 refers to the single node experiment in the caption. Furthermore, the number of nodes in the multiple node experiment only appears on Figure 8. Can it be a bit more prominent?
The the caption of Table 3 makes explicit that it refers to a single node experiment. The number of nodes appears at beginning of section 3.2.2 "Multiple nodes, multiple cores", in the first phrase: "We performed a larger-scale experiment at the Abel computer cluster1, using up to 256 cores on 16 nodes with Intel Xeon E5-2670 CPUs."

2.1.19

19. Why is S>2, S/n >1 for n=2 in table 4. Whilst some timing variation might be expected (e.g., perhaps due to the computer doing something else), one would still expect S to be some way short of 2.
It is true that Amdahl's law predicts S to be smaller than 2. However, in certain cases the parallel efficiency may be greater than 1 because of superlinear speedup. Superlinearity can be a consequence of memory usage (cache effect, changes in pagination) and storage. Another possible source are changes in network latency. We added a clarifying footnote.

2.1.20

20. Footnote 5 page 31: 'instructions' is misspelt.
The spelling was corrected.

2.1.21

21. For the layman, how does one set the number of nodes and cores (section 4.3)?
On a purely operational level, in order to set the number of cores, one would add an option to the command mpirun.
For example, one would type
`$ mpirun -n4 FLEXPART_MPI`
in order to run `FLEXPART_MPI` using 4 cores.
In practice the 'best' number of cores for a given simulation will depend for example on the size of the problem and the hardware availability. Choosing the 'right' number of cores may require some hands on experience. A comment was added.

2.1.22

22. Table 7: what are the units of MASS?
The units are described in table 10. These can be mass (e.g. kg) or 1 (unit-less as in the case of mixing ratio and backward mode). A clarification was added in the table.

2.1.23

23. I find Table 8 difficult to follow:

a. Earlier square brackets are used for the default value. Here they are used for the default unit and ordinary brackets for the default type. Furthermore on first reading, it is not clear that the word ?default? in the final column header refers to 'unit' and to 'type'. It might be clearer if it is written {Default unit} (Default type)? (i.e., putting default within the brackets and using a different type of brackets).

We have split the column in order to clarify the information. There is a column for type and another for the unit. There is not such a a thing as a default value for the species. The user has to make a choice of the the species to be transported. A common test can be carried out with species 24 (dry air).

b. Preldiff: What is the component D_i? Is it the diffusivity of the specie in the `SPECIES_nnn` file?

Yes, it is the diffusivity of the specie, that was clarified in the table.

c I note that the emission variation factors are specified in local time. Presumably, the time zone is given somewhere (or calculated based on location) and it must account for daylight saving time?

The astronomical local time is automatically calculated based on longitude, although there is no correction for Daylight Saving Time. Clarified in the Table.

2.1.24

24. The prepending p is missing from pdryvel and pdquer (lines 2 and 3, page 38).
The prepending p was added in both cases.

2.1.25

25. I don't think it is clear what 'up to 1 hour' (line 26, page 39) means. Does it mean that delta t >= 1 hr or delta t <= 1 hr? Similarly does 'up to 0.75 deg x 0.75 deg latitude / longitude resolution' (lines 28-29, page 39) mean delta lat >= 0.75 deg or delta lat >= 0.75 deg?
We have clarified the statements

2.1.26

26. Should `IND_RECEPTOR=4` for `grid_wetdep_date_nnn.nc`? (Note that Table 6 only has options 1 and 2 for IND_RECEPTOR - options 3 and 4 are missing.) Should IOUT =2, 3, 10, 11 (rather than 2, 3, 9, 11) for receptor_pptv?
For `grid_wetdep_date_nnn.nc`, the value of `IND_RECEPTOR` should be 4, the typo in Table 11 (former 9) was corrected. Options 3 and 4 were added to Table 6.
For `receptor_pptv`, IOUT=2, 3, 10, 11. This was also corrected in Table 11 (former 9).

2.1.27

27. Line 1 page 44: Should this read 'and nnn is the species number'? rather than 'and species is the species number'?
The typo in was corrected.

2.1.28

28. There are a number of occasions where the text runs off the page (e.g., line 15, page 58).
This particular case where an overfull occurs was fixed. The typeset version should eliminate all overfulls.

2.1.29

29. Is there a typo on line 8 page 45: 'verb|make|' ?
The typo was corrected.

2.1.30

30. Line 6, page 47: FLEXPART should be in capital letters.
Instead of FLEXPART we wrote the web address "flexpart.eu"

2.1.31

31. Table 12: Hypothetical is misspelt. Also, should Test name 12 under gridded output have a description of 'concentration and trajectory cluster in NetCDF'.

The typo was in 'Hypothetical' was corrected and removed the line in the table.

2.1.32

32. Line 1 page 50: Remove 'and' to read results from http://flexpart.eu?.

The word 'and' was removed. The user can obtain these reference results from https://flexpart.eu/

2.1.33

33. I think a minus sign is missing from the options of tar in a number of places (e.g., tar xf ... (line 19, page 53) should be tar -xf... Also lines 1, 7, 14, 21 (page 54).

The minus sign was added to the command line statements for `tar`.

2.1.34

34. As the authors state, GRIB-API support has been discontinued and ecCodes is the primary GRIB package at ECMWF. Appendix B1 appears to give instructions using `grib_api` only. Can these be updated to use ecCodes? Also, the meteorological files generated here (EA170122??) do not have a century label. How can pre 2000 met data files be generated and used?

We added a reference to `eccodes` in Appendix B1. The migration to eccodes is not yet complete in all FLEXPART versions. Legacy versions of FLEXPART run with grib api (and gribex). Despite substantial updates, flex_extract still contains a substantial amount of legacy code. We keep the grib_api instructions for backward consistency. We have consulted the FLEXPART community and some users appreciate the backward compatibility and the possibility of update without the need for installing new libraries. This is the case in particular for those users without IT support that need to manage their own installation. This backward compatibility does not affect the performance of FLAXPART because all required reading functionalities are currently supported by `grib_api`. Future FLEXPART and flex_extract versions will complete the migration and deprecate `grib_api`.

Pre-2000 met data files can be generated and used. In most cases there would not be overlapping decades. If the user prefers or needs to use full four digits years, it is a matter of modifying the string in flex_extract and adapt the AVAILABLE file accordingly. The length of the string containing the meteorological file names in the routine `readavail.f90` foresees 255 chars for the filename, that will usually suffice, unless a long path is included. Note that `mkavail.py` obtains the time from the GRIB information and not the filename, thus confusion is excluded safely. In `flex_extract` (i.e. in the extraction routines) these two digits can be added to the prefix. The next version of `flex_extract` will have standard four digit years.

2.1.35

35. Appendices C1 and C4: What is the time resolution of the meteorological data? The authors only give the spatial resolution.

The standard time resolution of the AROME meteorological data is 1 hour. The standard time resolution of the NorESM meteorological data is 3 hours. This was clarified in the text.

2.1.36

36. Line 8 page 61: Assuming this is the Thomson mentioned on lines 1-2 (Thomson, 1987), there is no 'p' in his surname.

The typo was corrected.

---

3 References

Anttila, T., Brus, D., Jaatinen, A., Hyvrinen, A.-P., Kiveks, N., Romakkaniemi, S., Komppula, M., and Lihavainen, H.: Relationships between particles, cloud condensation nuclei and cloud droplet activation during the third Pallas Cloud Experiment, Atmos. Chem. Phys., 12, 11435-11450, https://doi.org/10.5194/acp-12-11435-2012, 2012.

Hittmeir, S., Philipp, A., and Seibert, P.: A conservative reconstruction scheme for the interpolation of extensive quantities in the Lagrangian particle dispersion model FLEXPART, Geosci. Model Dev., 11, 2503-2523, https://doi.org/10.5194/gmd-11-2503-2018, 2018.

Pisso, I., Haynes, P. H., and Law, K. S.: Emission location dependent ozone depletion potentials for very short-lived halogenated species. Atmospheric Chemistry and Physics, 10(24), 12025?12036. http:doi.org10.5194acp-10-12025-2010, 2010.

Ristov, S, R Prodan, M Gusev and K Skala: Superlinear speedup in HPC systems: Why and when?, 2016 Federated Conference on Computer Science and Information Systems (FedCSIS), IEEE, 2016.

Stohl. A.: Characteristics of atmospheric transport into the Arctic troposphere, J. Geophys. Res., 111, D11306, doi:10.1029/2005JD006888, 2006.