

RDF Planner (v2.6): Users Guide

John Methven

07/03/2005

Department of Meteorology,
University of Reading,
PO Box 243, Earley Gate,
Reading, RG6 6BB.

J.Methven@reading.ac.uk

1 Visualising Domain Filling Trajectories

This document describes the features of an IDL program, RDFplanner, which can be used interactively to view particle trajectories released from a 3D grid and to print them. It is also possible to visualise emissions tracers on the grid (emissions accumulated along the complete length of each trajectory), the closest approach of trajectories to selected “aircraft bases” and the time at which this occurs. Finally, it is possible to view reverse domain filling reconstructions of chemical fields on the 3D grid (RDF3D).

The program is available as an IDL executable *rdfpplanner.sav* or as the source code *rdfpplanner.pro*.

The executable version of RDFplanner can be run using **IDL Virtual Machine** on a PC, MAC or UNIX workstation. The program will automatically detect which type of computer you are using. You do not need an IDL licence and can download IDL Virtual Machine for free from

<http://www.rsinc.com/download/index.asp>

The source code for the IDL program is not as easy to use as the executable but can be modified if desired (see Section 3). You need a full IDL licence to run the program. You will also need *rdftread_vm.pro* which contains the source code for reading trajectory files output by the UGAMP trajectory model (Methven 1997).

1.1 Domain filling trajectory data (utraj-df files)

The term “particle trajectory” describes the path of a point which is blown by a time-dependent wind field (i.e. (u, v, w) as a function of (x, y, z, t)). Trajectories following the analysed wind field are described by their coordinates (e.g. longitude, latitude, pressure) at regularly spaced time intervals. “Domain filling” refers to calculations where the arrival points of back trajectories (or release points for forward trajectories) form a dense, regular grid in a specified volume. The term “reverse” is used to indicate that the particles are followed backwards in time. Back trajectories are assumed to arrive on a 3D grid consisting of a stack of horizontal grids (regular in longitude and latitude) on a range of pressure levels. Forward trajectories are assumed to depart from similar grids. The trajectory length (time before arrival for back trajectories) is denoted by the letter T . Other fields can also be recorded following the trajectories: for example, temperature, specific humidity or potential vorticity. These extra fields are described as “attributes” and will be denoted by the variable C . The change in the value of an attribute over the length of a trajectory is denoted by $C(0) - C(T)$.

Note that because these files contain a 4D dataset (the trajectories arrive on a 3D grid and each one is described by a time series for its coordinates and attributes), their filesize can be very large. Therefore they are slow to download from the ITOP-UK forecast website.

However, RDFplanner can be used to view emissions tracers and chemical attributes (much smaller files because they are only 3D datasets), without the trajectory data itself, provided that plots of the associated trajectories are not required.

1.2 Emissions tracers, closest approach diagnostics and Lagrangian opportunities (utraj-em files)

“Emissions tracers” are calculated by integrating surface emissions along each trajectory when it is within the boundary layer. The surface emissions are specified using an inventory. During the ITOP campaign the EDGAR inventories for NO_x and isoprene emissions are used to indicate trajectories that are likely to have been influenced by anthropogenic or biogenic emissions respectively. The emissions from the surface are assumed to be instantaneously mixed throughout the boundary layer column so that they are equivalent to a volume source within the boundary layer. The boundary layer depth (time and space dependent) is obtained from the same numerical weather prediction model as provides the wind and temperature fields (usually the ECMWF model). Chemistry and dilution by mixing is not modelled along the trajectories, so the values assigned to back trajectories are not intended to represent concentrations on the arrival grid. Both NO_x and isoprene have short photochemical lifetimes compared with the length of trajectories used.

For each domain during the ICARTT campaign (US East Coast, Azores and W. Europe) emissions are accumulated along back trajectories from the 3D grids. In RDFplanner these results are referred to as “accumulated emissions”. In addition, forward trajectories are calculated from the same grids and emissions integrated along them are referred to as “future emissions”. When viewing utraj-em files associated with forward trajectories it is possible to plot both accumulated and future emissions on the 3D grid.

“Interception points” are defined as the closest approach of each trajectory to the chosen aircraft bases (for ICARTT these are: 1 - Pease, NH; 2 - Horta, Azores; 3 - Creil, France). You can view the great circle distance between the interception points and the corresponding base and also the time along the trajectory at which the closest approach occurs.

“Lagrangian opportunities” are defined by the following criteria:

1. Closest approach is less than R , the aircraft operational range from each base, for at least two bases. If the trajectory passes within range of 3 bases it is referred to as a “3-point opportunity” (otherwise 2-point).
2. Accumulated emissions (along back trajectories from the 3D grid) exceed Q , a chosen threshold value.
3. Future emissions (forward trajectories only) are less than Q in order to exclude trajectories that are anticipated to pick up more emissions from the boundary layer (and therefore loose identity between aircraft interceptions).

During the ICARTT forecasting period $R=1000\text{km}$ for all 3 bases and the emissions criterion is based on the NO_x tracer only with $Q=10\text{ppbv}$.

1.3 Chemical attributes from TOMCAT (tomcat-attr files)

“Chemical attributes” are found by interpolating chemical distributions (in space and time) from a global chemical transport model to the origin of each trajectory (using its full length). During the ICARTT campaign the TOMCAT global CTM is being run in near-real time (about 19 hours behind present) driven by wind analyses from the ECMWF. The back trajectories are sufficiently long that a TOMCAT chemical analysis exists even at the origin of forecast trajectories. For example, the longest forecast lead time for the Azores domain is 5 days but the back trajectories are 7 days long so that the TOMCAT fields dating from 2 days before the latest meteorological analysis are used to find the attributes. For the US East Coast domain the back trajectories are shorter (3 days long) but the longest lead time is also 3 days so that the chemical attributes can be calculated as soon as TOMCAT has been brought up to date with the latest ECMWF analyses.

2 Forecast dates and lead times: filename structure

The “base date” refers to the latest meteorological analysis from which the met forecast started.

The “verification date” refers to the time that a particular forecast represents the atmosphere. In the case of DF3D trajectories it is the time that the calculation is initialised on the 3D grid. For forward trajectories it is also referred to as the “release time”. For back trajectories it can be referred to as the “arrival time” since that is when air parcels associated with the back trajectories would reach the 3D grid.

The “forecast lead time” is the time elapsed between the base date and the verification date. In order to integrate trajectories forwards in time from the verification date, met forecasts for even longer lead times are required. Therefore forward trajectory forecasts are inherently less certain than back trajectory forecasts which use met forecasts with progressively shorter lead times, and then analyses, as trajectory length increases.

RDFplanner expects the data to have the following filename structure:

{FP}_{REG}{YYYYMMDDHH}{FFF} where

FP = forecast product (utraj-df, utraj-em or tomcat-attr)

REG = region of domain and indication of forward (f) or back (r) trajectories (usf, usr, azf, azr, eur)

YYYYMMDDHH = verification date in UTC (year, month, day, hour)

FFF = meteorological forecast lead time at the verification date (in hours).

3 Default setup and possible changes

3.1 Pre-compiled Executable

If you are using the pre-compiled executable *rdfplanner.sav* it is not possible to change the setup. Simply load *rdfplanner.sav* into IDL Virtual Machine and it will run. Use the file chooser that appears to select a data file with the verification date and lead time that you are interested in. The program will then read the appropriate utraj-df, utraj-em and tomcat-attr files if they are available in the same directory. Note that the program will function provided that at least one of the appropriate utraj-df or utraj-em files is available. This means that if your internet connection is slow there is no need to download the large utraj-df files to view emissions tracers and chemical attributes. Note that it will not work if only the tomcat-attr file is available.

3.2 Running the source code using IDL

These instructions refer to those people running *rdfplanner.pro* with a full IDL licence. The switches listed below can be found near the top of the main program (where the comments say MAIN PROGRAM after all the subroutines in *rdfplanner.pro*).

When you compile and run the code in the IDL environment you will be prompted to enter the verification date (YYYYMMDDHH), forecast lead time (in hours) and regional code (REG).

The wind data code (*IDATA*) is then used to construct the full path and filename for the input data. The default setting *IDATA=1* means that the data files are expected in the IDL default directory. Under UNIX this corresponds with your current directory. The simplest approach is to copy the data files for trajectories, emissions tracers and chemical attributes into the same directory and start IDL from that directory, using the default setting (*IDATA=1*). Under WINDOWS the IDL directory (IDL60 for version 6.0) is the default. You must put your data files into this directory to read them with *IDATA=1*.

If the data files are not in the default directory then you must edit *rdftread_vm.pro* and the subroutines RDFCHEM and READINPOS (in *rdfplanner.pro*) and add an appropriate file path in the CASE statements for *IDATA* (follow the existing examples).

Note that once you have selected the options for operational use you can by-pass most questions by editing the main program and setting *LQUEST=0*. Default values will then be used for *colfile*, *IDATA*, *ztype*, *t2th* and *logq*. These defaults can be changed by editing their values at the top of the main program.

IDATA : as above

ztype : 0 for flight level (hft), 1 for km

t2th : 0 - temperature; 1 - potential temperature; 2 - equivalent PT; 3 - wet bulb PT

logq : y - show log(specific humidity); n - show specific humidity

LTRAJ : Read trajectory data (utraj-df file)? - y or n

LEM : Read emissions tracer data (utraj-em file)? - y or n

LCHEM : Read chemical attribute data (tomcat-attr file)? - y or n

As explained for the executable, the program will function provided that either the utraj-df or utraj-em file is available (or both).

Only one thermodynamic variable and one moisture variable can be viewed and you must choose at this stage. The input data are usually temperature and specific humidity and the other options are calculated from these and pressure.

All flight track data must be in NASA Ames format with a filename of the form $\{FT\}_{\{PLANE\}}_{\{YYYYMMDD\}}_{R0}.ict$.

FT = flight track type (*utraj-fgr* for ground relative coordinates or *utraj-far* for coordinates relative to air-masses at the verification time for the DF3D trajectories).

PLANE = aircraft name (faam = FAAM BAe146, dc8 = NASA DC8, p3 = NOAA WP3D, fa = DLR Falcon flight A, fb = DLR Falcon flight B).

YYYYMMDD = verification date in UTC (year, month, day).

Again the file path is determined by the code *IDATA* and new filepaths can be set by editing the CASE statement for *IDATA* in the subroutine READINPOS of “rdfplanner.pro”.

Overlaying flight tracks has been disabled in the executable version *rdfplanner.sav*.

You can also edit the default values used for contouring each type of attribute and the default map range for the trajectory plots.

4 Using the plotting window

Once the input files have been selected, for the appropriate verification date and forecast lead time, the data will be read. Note that it may take some time to read in the trajectory data.

A window called “RDF planner” should then appear. You will see a number of buttons (IDL widgets) that control the plotting procedure and a draw window in which the plots appear. The bottom row of buttons enable plotting, printing and exit from RDFplanner. The function for each button is summarised in the table below.

Bottom row

- Your first plot must be a map of attribute values on a constant pressure surface of the arrival grid. This is obtained by clicking on the **plan** button.
- A cross-section through the arrival grid can only be examined after specifying the start and end coordinates of the cross-section. This is achieved by pressing the left mouse button at your chosen start point on the plan in the draw window (a cross will appear), holding down the mouse button while moving the mouse cursor to your chosen end point and releasing the mouse button there. The line of the section will be drawn across the plan. Each section follows the great circle between the start and end point.
- To plot the cross-section that you have just defined, click on the **section** button.
- You can view the trajectories arriving in the neighbourhood of a point of interest by clicking the left mouse button at any point on the cross-section in the draw window (a cross will appear) and then clicking on the **traj** button. The trajectory arriving closest to the point of interest is the thick line and the other six trajectories are the neighbouring trajectories on the arrival grid to the North, South, East, West, above and below. The trajectories are coloured by the value of the attribute $C(T)$ so that the colour of the thick trajectory matches the colour (using colour fill) on the section and plan at the trajectory point of interest marked by the cross.
- You can now view the plan, cross-section or trajectories at will by clicking on the appropriate button.
- Note that you cannot select the “trajectory point of interest” by clicking on a plan.
- **Image** outputs the current picture to a PNG image file in the current directory (UNIX) or the default IDL directory (WINDOWS).
- **Print** outputs a postscript file of the current picture to the current directory (UNIX) or sends the image directly to your PC’s default printer (WINDOWS).
- **Verbose** toggles between enabling or disabling (limited) help messages.
- **About** gives author details.
- To exit RDFplanner click on the **done** button.

Top Row

- Pick from a drop-down list of trajectory “attributes”.
- Show the value of an attribute interpolated from ECMWF model fields (analyses or forecasts) to the chosen time-point (T) along every trajectory or show the change in attribute between that point and the verification date. It is only possible to show the change in attribute for data from the trajectory file (utraj-df) and not for the chemical attributes. Accumulated emissions, future emissions, intercept distance/time and chemical attributes are unaffected by this button.

- Choose the time-point along the trajectory used to select the attribute value. This also selects the length of trajectory (T) shown on the trajectory plots. Accumulated emissions, future emissions, intercept distance/time and chemical attributes are unaffected by this button and if they are chosen as the attribute the time-point will automatically be reset to the end of the trajectory (i.e., its maximum length). If you have opted not to read the trajectory data ($LTRAJ='n'$) then the trajectory length will be labelled as zero on the button, although obviously the diagnostics have been calculated with trajectories of non-zero length.

Second Row

- The “pressure”, “up” and “down” buttons change the level used for the plan, the level of the trajectory point of interest and the position of the corresponding cross on the section.
- The default vertical axis for cross-sections is pressure but it is possible to toggle to a height coordinate using the “p or z axis” button (see table). *ztype* (see above) determines whether the height units are km or flight level (FL in 100× feet). Note that these coordinates are more precisely “pseudo-height” and are functions of pressure only, assuming hydrostatic balance and a standard atmospheric profile (as used in the aviation industry). No account is taken of variations in surface pressure or the vertical temperature profile. The conversion formula from pressure (hPa) to pseudo-height (km) is $z = 44.33077 * (1. - (p/1013.25)^{0.19026})$.
- You can change the ranges of the plots in longitude and latitude. Different ranges are used for plans and trajectory plots. The selected ranges are remembered if you switch between plans and trajectory plots.

Third Row

- You can change the minimum attribute value used for contouring and the contour interval (dC - must be positive).
- Choose between line contours, colour fill or both together.
- Select a different colour scale.
- Reverse the colour scale.

5 Interpretation of Results

The trajectories are calculated using the 3D wind field from atmospheric analyses or forecasts. Details of the method used by the UGAMP trajectory code can be found in Methven (1997). The result of trajectory calculations are the coordinates (longitude, latitude, pressure) along

the trajectories as a function of time. All other attributes (like specific humidity) are obtained by interpolating the analysed fields in time and space to every point along the trajectories. No explicit calculation of the change in specific humidity (or any other attribute) following an air-mass is made. This can make interpretation of the fields plotted by RDFplanner rather confusing. The following notes focussing on back trajectories may aid interpretation. For more detail and examples using RDFplanner to analyse flight data refer to Methven *et al.* (2003) and Purvis *et al.* (2003).

1. A plot of an attribute's value at $T = 0$ on the arrival grid is just a plot of the field from the weather centre's model at the time of arrival (analysis or forecast). Typically, the trajectory arrival grid has a much finer spacing than the analysis grid from the weather centre (e.g. ECMWF) so that the $T = 0$ fields look very smooth.
2. A plot of an attribute's value for other trajectory lengths ($T > 0$) can be interpreted as the view of the field that would have been realised if that attribute had not changed along the length of the trajectories. In other words if that attribute was materially conserved. This is a sensible interpretation for long-lived tracers. For example, if the analysed flow was adiabatic and frictionless then PV would be conserved and the RDF picture would be a good representation of the real PV structure in the atmosphere. There is usually much more structure in the RDF picture of PV than in the analysis because the PV analysis is limited in resolution. Structure will appear on ever finer scales in the RDF picture as T is increased. A sensible choice for T must be made so that the RDF picture does not become too noisy (because the arrival grid spacing is too wide compared with the tracer scales it is trying to represent). Experience suggests that when the arrival grid spacing is 25km an optimum trajectory length is about 1.5-4 days depending on the rate of separation of trajectories (Methven *et al.* (2003)).

However, the full back trajectory length is always used for chemical attributes, emissions tracers and the Lagrangian opportunity diagnostics. As a result these attributes are often quite noisy (especially on the European domain because it is furthest from the North American emissions and the polluted air masses have had many days to stretch and thin).

3. It is not appropriate to interpret in this way the plots for attributes that are not conserved, such as pressure. Indeed, it is better to plot the change in these attributes over the trajectory as a function of their location on the arrival grid. For example, the change in pressure along trajectories immediately highlights regions where air has experienced ascent or descent during the period before arrival. The change in potential temperature highlights air-masses that have experienced heating or cooling. For these kinds of attributes $C(T)$ can be used as an indicator of the origin of the trajectory even if attribute C undergoes major modification along trajectories. For example latitude(T) or $p(T)$ can be used to indicate the origin of particular air-masses. Specific humidity $q(T)$ can be used to indicate air-masses that were very dry at time $t_0 - T$ (stratospheric in origin) or very moist (possibly in the boundary layer).

4. Interpretation of the change in attribute for long-lived tracers (like PV) is difficult. It appears that the value of analysed PV following trajectories mostly changes due to the relatively coarse resolution of the analyses and the large numerical dissipation that this implies. The correspondence of the $T = 3$ day RDF3D PV picture with UKMO mesoscale model results and ozone data collected during the NERC DCFZ campaign, indicates that this bears a much closer resemblance to the atmospheric PV structure than the ECMWF analyses themselves. In other words, PV is more nearly conserved in the atmosphere than in the analyses. Sections depicting the change in PV along trajectories show that during the last day before arrival the finescale PV structure is dissipated rapidly in the analyses so that tropopause folds do not appear to intrude as deep into the troposphere as is observed. It is possible to use the RDF3D picture as a good indicator of the finescale structure in PV and other long-lived tracers like CO or ozone.
5. Interpretation of the specific humidity attribute is especially difficult. In unsaturated air-masses changes in analysed q along trajectories occur due to the coarse resolution of the analyses for the same reason as for PV. However, q may change by several orders of magnitude following saturated air-masses due to condensation. You can also view equivalent potential temperature, θ_e , or wet-bulb potential temperature, θ_w , along trajectories. These are conserved following unsaturated or saturated air-masses in the absence of sensible heating. If θ_e is nearly conserved for an air-mass then it is possible to use changes in specific humidity as an indication of condensation or sub-grid scale humidity fluxes into the air-mass.

References

- Methven, J. (1997). Offline trajectories: Calculation and accuracy. Technical Report 44, U.K. Univ. Global Atmos. Modelling Prog., Dept. of Meteorol., Univ. of Reading, U.K.
- Methven, J., Arnold, S., O'Connor, F., Barjat, H., Dewey, K., Kent, J., and Brough, N. (2003). Estimating photochemically produced ozone throughout a domain using flight data and a Lagrangian model. *J. Geophys. Res.*, **108(D9)**, 4271, doi:10.1029/2002JD002955.
- Purvis, R., Lewis, A., Carney, R., McQuaid, J., Arnold, S., Methven, J., Barjat, H., Dewey, K., Kent, J., Monks, P., Carpenter, L., Brough, N., Penkett, S., and Reeves, C. (2003). Rapid uplift of non-methane hydrocarbons in a cold front over central Europe. *J. Geophys. Res.*, **108(D7)**, 4224, doi:10.1029/2002JD002521.

Widget name	Position	Function
Attribute	top row	Choose the attribute C to be displayed.
Type	top row	Choose to view the <i>attribute value</i> $C(T)$ or the <i>change in attribute</i> along each trajectory $C(0) - C(T)$. On plans and sections the chosen attribute type labels each trajectory and is displayed on the arrival grid as a contour plot. Note that only the <i>attribute value</i> can be shown on trajectory plots.
T	top row	Select the trajectory length (in days).
Pressure	row 2	Select the pressure level for plotting with the plan button. Also moves “trajectory point of interest” (shown as a cross on sections).
Up	row 2	Move up an altitude level (to lower pressure).
Down	row 2	Move down an altitude level (to higher pressure).
p or z axis	row 2	Toggle between pressure or pressure-height for vertical axis on sections. Pressure-height will be shown in km or as flight level (100× feet).
min(lat)	row 2	Southern edge for plan or trajectory plot (degrees). Edit to change.
min(lon)	row 2	Western edge longitude.
max(lat)	row 2	Northern edge latitude.
max(lon)	row 2	Eastern edge longitude.
min(C)	row 3	The first contour level. Edit to change its value.
dC	row 3	The contour interval. Edit to change.
Contour style	row 3	Choose labelled contours, colour fill or both together.
Colour table	row 3	Choose from the standard IDL colour tables.
Reverse	row 3	Reverse the colour table.
Plan	row 4	View the chosen attribute on a pressure surface from the arrival grid.
Section	row 4	View the chosen attribute on the cross-section defined using the mouse dragging procedure described above.
Traj	row 4	View the trajectories arriving in the neighbourhood of a point on the section preselected using the mouse.
Image	row 4	Output the current plot as a PNG image file in the default IDL directory.
Print	row 4	PCs: Output the current plot directly to printer. Workstations: Output as a postscript file in the current directory.
Verbose	row 4	Toggle to switch help messages and extra text on or off.
About	row 4	Author information.
Done	row 4	Exit RDFplanner.

Table 1: The function of each widget in the RDFplanner window