August 1996 ADMS 2
P13/03C/96

AVERAGING TIME AND FLUCTUATIONS IN ADMS VERSIONS 1 AND 2

by
David J. Thomson, Meteorological Office

Summary This document describes the treatment of averaging time and release duration and the way output from the fluctuations module should be interpreted in ADMS versions 1 and 2.

---

1 Introduction

The treatment of averaging times in ADMS is quite complex, especially in connection with the treatment of σ_y and of fluctuations. The purpose of this note is to clarify the situation on ADMS version 1, to relate it to the approach used in the R91 report (Clark et al 1979), and to discuss the changes for version 2.

2 Averaging time and ensemble averages

No dispersion model attempts to predict *the* concentration which occurs at a particular point and time of interest. For identical gross external conditions (e.g. for fixed geostrophic wind, surface heat flux, boundary layer depth and location), the results will vary due to the unpredictable nature of the turbulence. As a result, models aim to predict some sort of average concentration. Some models also aim to predict the concentration probability distribution.

The simplest type of average conceptually is the ensemble average. One considers an ensemble of realisations of the flow in which the gross external conditions are identical but in which the details of the turbulence differ. This is especially straightforward in a wind tunnel where one can envisage repeating an experiment many times. However it is not so straightforward in the atmosphere for two reasons. Firstly, it's not so easy to say what the 'gross external conditions' are and distinguish them cleanly from the 'turbulence'. If the geostrophic wind, surface heat flux, boundary layer depth and terrain are uniform in space and time over the domain of interest then one can define the gross external conditions as consisting of these aspects of the situation. More generally however there are mesoscale flows and no clear spectral gap which complicates the situation. Secondly one cannot of course repeat an atmospheric dispersion scenario many times with identical external conditions.

As an alternative one can consider time averages (at least for continuous sources). Again this is straightforward in a wind tunnel but not so straightforward in the atmosphere because the external conditions change with time – if one averages for so long as to smooth out all variability due to the turbulence then, except in rather special condition, the meteorology will have changed considerably. In general there is no averaging time, which removes all the unpredictable fluctuations (except perhaps very long times – climatological averages – for continuous sources). Time averages over periods of various lengths may also be of interest for their own sake and not just for the purpose of smoothing out the unpredictable fluctuations in concentration.

We have seen that time averaging cannot remove all the unpredictable fluctuations and so models which give single values for concentration must be considering, implicitly or explicitly, some sort of ensemble averaging too. There are (at least) two possibilities here – one can define an average over an ensemble E1 of different realisations of the turbulence with specified meteorology which is constant in space and time (in so far as one can define the differences between turbulence and changes in meteorology) or one can average over an ensemble E2 of realisations which are consistent with the met observations (see discussion in Weil et al 1992). The second choice implies (conceptually at least) averaging over that sub-ensemble of the climatological ensemble which consists of all occasions which are consistent with the met observations. This may include realisations in which the met as well as the turbulence evolves differently (depending on how tightly the observations constrain the meteorology). The first is perhaps the easiest conceptually while the second may be the more useful. The second doesn't require any conceptual distinction between turbulence and changes in meteorology, although such distinctions are likely to be needed in constructing models.

What does this mean in practice? For simplicity we consider only the continuous source case. Consider first a case where we have a single set of met observations (e.g. geostophic wind, surface heat flux and boundary layer depth, or 10m wind , cloud cover, time of day and time of year – it is convenient to regard time of day and year as part of the met) and consider the ensemble E1 (see figure 1 (a)). Because we are averaging over the turbulence and assuming constant meteorology there is no change in the plume width or concentration with averaging time. Now consider the ensemble E2, again with a single set of met data. For small averaging times this will give similar results since the assumption of constant met is then a good approximation. As the averaging time increases, changes in meteorology and, in particular, consequent changes in wind direction become more important, resulting in a wider plume (figure 1 (b)). This is often treated in models by simply increasing σ_y. If, still using E2, we had observation of meteorology at several times, this would give more information on the changes in meteorology and would result in a smaller ensemble (only those realisations consistent with the observed met would be retained) and a plume intermediate in width between that in figures 1(a) and (b). We could also consider the ensemble E2 in the case of no met measurements! The ensemble E2 consists of all occasions which are consistent with the observed met and so, in this case, constitutes the entire climatological ensemble. This corresponds to the 'long-term' model in R91 or to the long-term mean concentrations from ADMS statistical calculations. Finally one could consider other ensembles consisting of, e.g., all mornings in summer. This can be regarded as an E2 ensemble if we regard time of day and year as part of the met.

For a long averaging time with the E2 ensemble and a single set of met data, the results could depend on the time the met data is obtained (see figure 2(a) and (b)). It could be argued for example that, in estimating the area at risk from an accidental release with a wind direction observed at (or averaged over some period near) the start of the accident, one should use a larger σ_y than if one had a wind direction averaged over the release period or measured at (or averaged over a period near) the middle of the release. In a more extreme case one might have a wind observation which occurred sometime before (or after) the release.

A further possible complication is the use of alternative frames of reference in the way the different realisations are averaged, namely the aligning of the plume centroids in the cross-stream direction averaging. This results in narrower plumes and higher centre line concentrations (figure 3). (In mathematical terms, this can be understood as being due to the σ_y's and the maximum concentrations in the individual realisations having means which are not equal to the σ_y and maximum concentration of the mean concentration field.) This will give a better idea of the width and peak value in any individual realisation (since it avoids contributions to the plume width resulting from the scatter in the centroid positions), although the true peak may be considerably larger due to in-plume fluctuations (figure 3(b)). As an example, consider a continuous source and a short averaging time of a few minutes or less. The position of the plume averaged over such a short time will be unpredictable and will be in different positions in different realisations (see figure 3(a)). In this case the centroid aligning approach would predict a mean plume profile as in figure3(b) as compared with the non-centroid–aligned profile in figure 3(a). In this situation the approach is of somewhat doubtful value because (i) users may erroneously conclude that there is no risk at e.g. point P in figure 3(c), and (ii) the approach will still underestimate the peak concentrations because of the neglect of in-plume fluctuations and because of the lack of a similar treatment in the vertical (especially for elevated sources). Over longer periods in excess of the time scales associated with boundary layer turbulence, the scatter in the centroid positions is due mainly to differences in meteorology (assuming ensemble E2 with a single set of met data – with ensemble E1 the scatter between different realisations will become small). Here the approach has more to commend it. This is because the in-plume fluctuations and vertical meandering will be smoothed out to some extent by the large averaging time. Also the approach will reduce the significance of the problem discussed in the previous paragraph. However the approach will still underestimate the area at risk. To avoid treating short and long averaging times differently, one can consider (conceptually at least) averaging the concentration field over the turbulence (ensemble E1), then aligning the centroids, and finally averaging over the realisations of the meteorology. In the following 'centroid aligning' will always refer to this type of centroid aligning procedure.

The complications discussed in the previous two paragraphs don't alter the fact that σ_y tends to increase with averaging time. However they will alter the size of σ_y.

3 R91 and ADMS

We have seen that the averaging times and their interaction with ensemble averages is quite complex and subtle. It is therefore important to be clear what models are attempting to predict. ADMS predicts average concentrations over the ensemble E2 and it seems reasonable to assume this is also the case for R91. Also ADMS predicts the distribution of fluctuations about the mean in some situations. In ADMS and R91 each individual dispersion calculation uses a single set of met data. In the following we will always be concerned with the E2 ensembles defined by a single set of met data. Throughout the following t_{av} and t_R will denote the averaging period (i.e. the period over which any concentration measurement is made) and the release duration respectively.

In both R91 and ADMS the effect of changes in met are represented simply by enhancing σ_y. In both models this is done by enhancing the lateral spread due to turbulence σ_{yt} by a term σ_{yw} which represents changes in meteorological wind direction over a 'sampling time' t_s :

with $$\sigma_y^2 = \sigma_{yt}^2 + \sigma^2{}_{yw}$$

and $$\sigma_{yw} = \sigma_{\theta x}$$

$$\sigma_\theta = 0.065\sqrt{7t_s / u_{10}}$$

where u_{10} is the 10m wind speed and σ_θ is the standard deviation of the meteorological wind direction over the period t_s (the 'meteorological' in 'meteorological wind direction' indicates that we are considering changes in wind direction due to changes in met only, and we are not considering changes in wind direction due to turbulence). It is also possible to specify a value of σ_θ directly and thereby avoid using (1). t_s (or σ_θ if it is input) determines the extent to which the meteorological wind direction is assumed to change, and this corresponds conceptually to the extent to which the ensemble E2 is bigger than the ensemble E1.

3.1 The choice of t_s

The appropriate value to use for t_s and its relation to t_{av} and t_R is a little complex. We will discuss this first in the general context of models of the R91/ADMS type which use a 'sampling time' t_s to determine the σ_y enhancement. Subsequent sections discuss R91, ADMS 1 and ADMS 2 in more detail.

If we consider a single concentration measurement and a single source, and we also assume that the wind direction used in the dispersion model is the appropriate ‘mean meteorological wind direction for the period of interest’ (more precisely, the time mean of the meteorological wind direction over the period over which conditions affect the concentration which we are trying to estimate with the dispersion model) then is appropriate to take t_s equal to the length of the period of interest. This can be estimated as min(t_{av},t_R). This is not quite correct if t_{av} and t_R are both small compared to the travel time form the source to the receptor. However t_s is only used in assessing how much the meteorological wind direction might change over the period of interest. Hence the use of t_s = min (t_{av},t_R) is often justified because, in short range models, the met (and hence the meteorological wind direction) usually changes little over the travel time.

In other situations (i.e. if we are considering more than one concentration measurement or source, or if the wind direction used is not the appropriate mean meteorological wind direction for the period of interest, or if the travel time can’t be neglected) there may be reasons for making a different choice for t_s . Suppose first that the wind direction used is not the appropriate mean meteorological wind direction for the period of interest. In reality this is quite likely to be the case due to the siting of the wind vane, or due to the met observation not being averaged over the period of interest, or due to the fact that the meteorological wind direction experienced by the plume will be different from that at a fixed, point, or due to the inherent unpredictability of short time scale fluctuations in the meteorological wind direction. If the user is considering centroid aligning then this makes no difference and t_a= min (t_{av},t_R) is still appropriate. If no centroid aligning is considered these effects will all result in a wider ensemble average plume due to the uncertainty in the meteorological wind direction experienced by the plume. A possible approach is to impose a minimum value T_{min} of order 1hr on t_s. This leads to t_s = max (min (t_{av},t_R), T_{min}).

A second possibility is that the effective period of interest could be longer than min (t_{av} t_R) due to the travel time of the plume not being short compared to both t_{av} and t_R. This could also be treated by imposing a lower limit on t_s but, in contrast to the situation above, the effect will be present whether or not centroid aligning is used. This is because of along wind spread which means that different parts of the plume experience different meteorological wind directions resulting in an increase in the width of the plume in each realisation as well an increase in the width of the ensemble mean plume.

Other cases where one might wish to make a difference choice for t_s are if one is considering more than one concentration measurement. For example one might be considering many short period average values obtained from a continuous source over a longer period T or one might wish to consider a finite duration release where one is interested in the time evolution of the concentration at one or more points. In such cases (at least if T or t_R is long enough) the wind direction input into the model cannot be the appropriate mean meteorological wind direction for each of the concentration measurements. It is also unlikely to be equally representative for each concentration measurement. For example, if it’s a short period average meteorological wind direction, it will be most representative for the concentrations measurements occurring at that time with the (E2) ensemble mean plume width being larger at other times (see figure 4 (a)). Alternatively, if it’s a longer period average it will probably be most representative for measurements occurring near the centre of the period (see figure 4 (b)). If one is considering centriod aligning for each concentration measurement separately, then none of this matters and t_s= min (t_{zv},t_R) *or* t_s=max (min (t_{av},t_R),T_{min}) is still appropriate. If not however, this ensemble mean plume width will, in reality, vary in time as illustrated in figure 4. This cannot be treated easily in models of the R91 or ADMS type. However if we make the simplifying (but incorrect in detail) assumption that the input wind direction is equally representative for each concentration measurement, then the ensemble average plume width becomes constant in time. t_s should then be chosen to be T or t_R, possibly with a minimum value T_{min} imposed to account for the fact that the input wind direction may not be the appropriate mean meteorological wind direction for the wind period T or t_R in question or the fact that the travel time may not be short compared to T or t_R. In general for many measurements each of duration t_{av} over a period T (t_{av} « T), the appropriate value of t_s is (again assuming no centriod aligning) min (T, t_R) or max (min (T, t_R), T_{min}).

For multiple sources where the sources don't all have the same release period the choice is difficult. In reality, the E2 ensemble average concentration filed will have the width of each plume depending on t_{av}, t_R and on how close the release time and averaging time is to the time of the met measurement. As in the previous paragraph this can't be easily treated in R91 or ADMS type models and simplifications are necessary. Note that it doesn't in general make sense to relate t_s for each source to t_{av} and t_R only, nor does it make sense in general to consider centriod aligning for each source separately – consider for example a 2 hr release modelled as a single release and as two successive 1 hr releases. The best choice of t_s will probably need to be made on a case by case basis. A possible solution if one has a time series of met data – hourly say – is to consider the releases occurring in each hour separately (so that $t_R \leq 1$ hr for each source), use the appropriate met data for each release, and set t_s = 1hr to represent shorter time-scale variations in meteorology.

Finally we note that, for comparison with wind tunnel data, it is appropriate to set t_s (or σ_θ) to zero (ensemble E1).

The effect of the above choices on the long-term (climatological) statistics of concentration deserves a few comments. Such long-term statistics are generally calculated by calculating dispersion for a large number of met conditions representative of the climatology. The long-term mean concentrations should be insensitive to the choice of t_s and may be completely independent of t_s , depending on how the calculation is done). Long term percentiles of ensemble mean concentrations are however sensitive since the spatial peaks of the ensemble mean concentration fields are reduced if t_s is increased. If however probability distributions of concentrations are calculated for each set of met data (e.g. using the ADMS fluctuations module) and the long term frequency of exceeding a given concentration is calculated by combining these probabilities, then this should again be insensitive to t_s (for fixed t_{av} and t_R). This is because the reduced peak ensemble mean concentrations should be compensated by an increase in the predicted size of fluctuations.

In the following we discuss in more detail what has been done in R91 and ADMS with most of the discussion centring on the value of t_s.

3.2 R91

R91 only considers continuous releases (t_R infinite or much larger than t_{av} and travel time) or time-integrated concentrations from a finite duration release (t_{av} infinite or much larger than t_R and travel time). t_s is set equal to min (t_{av},t_R) with the caveat that the model should be used only when t_s is greater than about 30 minutes. The reason for this restriction is not entirely clear. It may be considered that, over shorter durations, the fluctuations will be large and so the actual plume will not approximate the Gaussain plume of R91. However this is probably also true (although to a lesser extent) over longer periods and in, any case R91 predictions are presumably intended to be interpreted as some sort of ensemble average. Alternatively it might be considered that, for shorter durations, the wind direction input to the model is unlikely to be the most appropriate or that the effective period of interest is longer than min (t_{av}, t_R) due to the travel time of the plume not being short compared to both t_{av} and t_R. In this case it could be argued that R91 could still be used but with the sampling time for changes in meteorological wind direction, t_s, set to 30 minutes if min (t_{av},t_R), is less then this (i.e. T_{min} is taken to be 30 minutes). Some users of R91 allow t_s to be less than 30 minutes. This may also be justified in some situations as discussed in § 3.1 above.

The view is taken that the turbulent spread corresponds to about a 3 minute release duration (or taken a three minute average concentration for a longer release) and that increases in σ_y over longer periods are due to changes in meteorology. This seems a little unsatisfactory, especially in convective conditions where turbulence time scales can been quite large (of order 15 minutes). In fact however, although '3 minutes' is mentioned in the description of the model, this value does not actually occur in the model itself and, in convective conditions, σ_{yw} does not exceed σ_{yt} for times up to about 15 minutes. In contrast there is an inconsistency in stable conditions where σ_{yw} for a 3 minute average exceeds σ_y, which is supposed to be a 3 minute average value.

3.3 ADMS

ADMS performs 3 types of calculations. The first calculates concentrations for a continuous release (t_R infinite or much greater than t_{av} and travel time), the second time integrated concentrations from a finite duration release (t_{av} infinite or much greater than t_R and travel time), and the third relates to instantaneous concentrations from a finite duration release (t_{av}= 0, $t_R$finite). The third type of calculation involves considering the along-wind spread and the 'tails' at the front and rear of the cloud of contaminant. Only the first two of these calculation types are dealt with by R91. We will now discuss each of these cases in relation to ADMS 1 and 2.

3.3.1 ADMS 1

In ADMS 1 it was decided to impose a lower limit of t_s of 1 hr. Conceptually this corresponds to setting T_{min} to 1 hr (see § 3.1 above). This leads to t_s = max (min(t_{av},t_R), 1 hr) although this has not been implemented consistently throughout.

It is useful to discuss this in more detail for each of the three types of calculation. In type 1 calculations σ_{yw} *is* evaluated with t_s equal to the 'averaging time' entered on the 'Mean Concentrations!' menu. This menu imposes a lower limit of 1hr on the value entered, consistent with the above philosophy. The value entered is used only for determining σ_{yw}. Although the value entered is called 'averaging time' it corresponds to what we have here called sampling time. For type 2 and 3 calculations σ_{yw} should be evaluated using the periods max (t_R, 1hr) and 1 hr respectively. In fact however (at least up to version 1.3) σ_{yw} is evaluated using a period of length zero (i.e. $\sigma_{yw} = 0$) for these cases. This is inconsistent with the general philosophy of ADMS 1 as outlined above and is hard to justify for the type 2 cases if t_R is non-zero. For the type 2 cases with $t_R = 0$ and the type 3 cases however, it is consistent with t_S = min (t_{av},t_R) which as we have seen may be appropriate for some calculations.

The fluctuation module predicts the distribution of values in the individual ensemble members around the mean value. The fluctuation module is only capable of calculating fluctuations due to the turbulence, although on a pragmatic basis the fluctuations module regards all eddies of time–scales less than 1 hr as turbulence (i.e. σ_{yw} for a value for t_{s} of 1 hr or less is regarded as part of the turbulent spread). Hence the module should only be used for cases in which σ_{yw} is evaluated using a value for t_{s} of 1 h for less, i.e. for type 1 calculations with t_{s}=1hr or for type 2 or for type 3 calculations (for which t_{s}=0). If t_{s} had been set to max (t_R, 1hr) for the type 2 cases as recommended above, then the fluctuations module would have been suitable for the type 2 case only if $t_R \leq$ 1hr.

For type 1 calculations, t_{av} affects only the fluctuations because t_s is specified directly by the user on the 'Mean Concentrations!' menu and is not calculated from t_{av}. Hence it was decided to input t_{av} on the fluctuations menu. It can be usefully thought of as 'the fluctuations averaging time'. If $t_{av} <<$ 1 hr one could imagine many measurements of concentration over successive periods each of length t_{av} and occupying 1 hr total (see discussion in § 3.1). If we assume that the wind direction input to the model is equally representative for each individual concentration measurement, then the distribution of such concentration measurements will approximate the distribution of any individual measurement over the ensemble and so the fluctuations module can be used to estimate such distributions.

Because t_{av} and t_s are specified independently in type 1 calculations they won't satisfy t_s = max (t_{av},1 hr) if t_{av} is greater than 1 hr (t_{s} always equals 1 hr for type 1 fluctuation calculations). Conceptually what happens here is that as t_{av} is increased the turbulent fluctuations are smoothed out but that, because t_s is fixed at 1 hr, the plume doesn't widen and fluctuations due to larger scale motions are not introduced.

3.2.2 ADMS 2

In ADMS 2 the approach is changed somewhat. The primary reason for the change is the dissatisfaction of users with the idea that σ_y is unchanged if the averaging or release time is reduced below 1 hr. The opportunity has also been taken to tidy up some of the inconsistencies noted above. The main change in approach is that t_s can now always be specified independently from t_{av} and t_R. t_s is specified on the 'Setup' menu. As in ADMS 1, σ_θ can still be specified directly if desired with the difference that σ_θ is now input with the other met parameters and so can be made to vary from one set of met data to the next. A second change is that the fluctuations module has been modified so that it can work for any value of t_s. The modifications are described in P13/04. As before the fluctuations module predicts the distribution of values in the individual ensemble members around the mean value. In allowing for any value of t_s, the module now accounts for fluctuations caused by both boundary layer turbulence and changes in meteorological wind direction. As a result the module no longer needs to take the pragmatic step of regarding all eddies of time scale less than 1 hour as boundary layer turbulence.

It remains to discuss the choice of t_s. The most common options are min (t_{av},t_R) and max(min(t_{av},t_R), 1hr). The 1 hr scale in max(min(t_{av},t_R), 1hr) could also be replaced by something different – 1hr corresponds to the approach taken in ADMS 1 while 30 minutes would correspond to R91. Reasons for these choices have been discussed in § 3.1 above. §3.1 discusses situations where one might with to choose a different value for t_s.

For type 1 cases the most common options are $t_s = t_{av}$ and $t_s =$ max (t_{av},1hr) and a default of $t_s =$ max (t_{av},1 hr) is recommended as appropriate for many calculations. For type 2 releases the most common options are $t_s = t_R$ and $t_s =$ max (t_R,1 hr) and a default of $t_s = t_R$ is recommended as appropriate for many calculations. Note however that this is not generally suitable for multiple source calculations where the sources don't all have the same release period (see discussion in §3.1). Finally, for the type 3 cases the most common options are $t_s = 0$ and $t_s =$ 1hr and a default of $t_s = 0$ is recommended as appropriate for many calculations. There is an asymmetry in these default recommendations which, by using a 1 hr mimimum for type 1 calculations only, favours a larger area at risk for type 1 calculations and a larger maximum concentration for type 2 and 3 calculations. Values of t_{av}, t_R, and t_s for the three types of calculations are summarised in the following table:

	Release duration t_R	Fluctuation averaging time t_{av}	Sampling time t_s
Type 1 'plume'	∞	t_{av}	t_s default: max (t_{av}, 1hr)
Type 2 'time integrated puff'	t_R	∞	t_s default: t_R
Type 3 'instantaneous puff'	t_R	0	t_s default:0

For type 1 calculations, t_{av} affects only the fluctuations as in ADMS 1 because t_s is specified directly by the user and is not calculated from t_{av}. Hence, again as in ADMS 1, it was decided to input t_{av} on the fluctuations menu. It can be usefully thought of as the 'fluctuations averaging time.' One could imagine many measurements of concentration over successive periods each of length t_{av} and occupying a period T in total (see discussion in § 3.1). If we assume that the wind direction input to the model is equally representative for each individual concentration measurement, then the distribution of such concentration measurements will approximate the distribution of any individual measurement over the ensemble and so, with the choice $t_s = T$ or $t_s = \max(T, 1\text{ hr})$, the fluctuations module can be used to estimate such distributions.

Because t_{av} and t_s (type 1 calculations) and t_R and t_s (type 2 calculations) are specified independently, it is possible to make t_{av} (type 1 calculations) or t_R (type 2 calculations) larger than t_s. Conceptually what happens here is that as t_{av} or t_R is increased the turbulent fluctuations are smoothed out but that, because t_s is fixed, the plume doesn't widen and fluctuations due to larger scale motions are not introduced.

References

Weil J.C., Sykes R.I. and Venkatram A. 1992 'Evaluating air-quality models: review and outlook' ,J.*Appl.Met.,*31, 1121-1145.

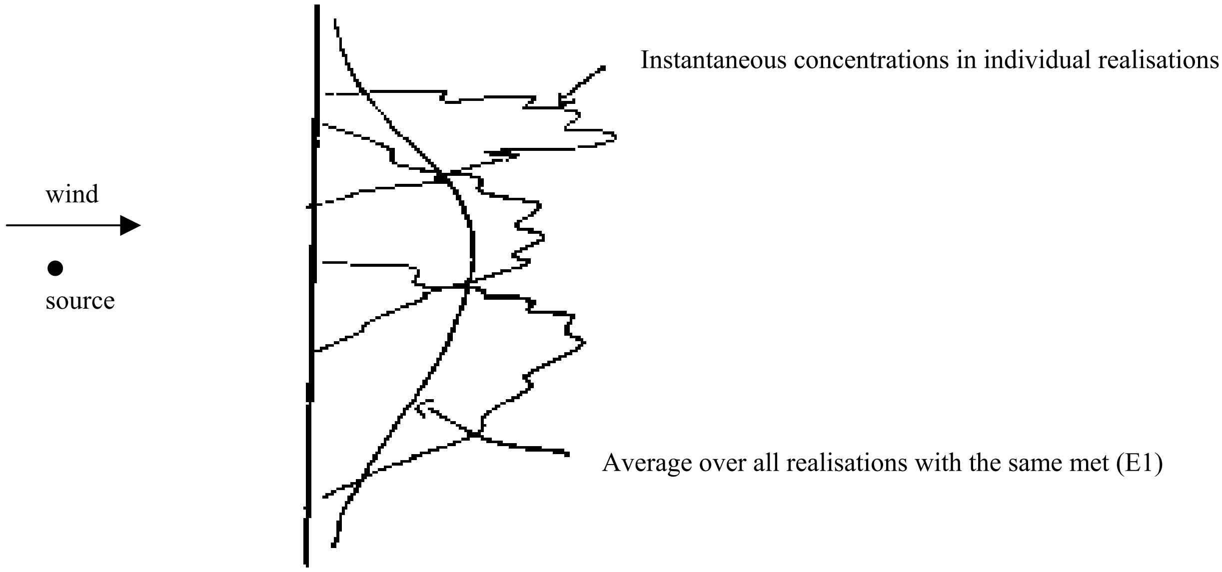

Figure 1a E1 averaging

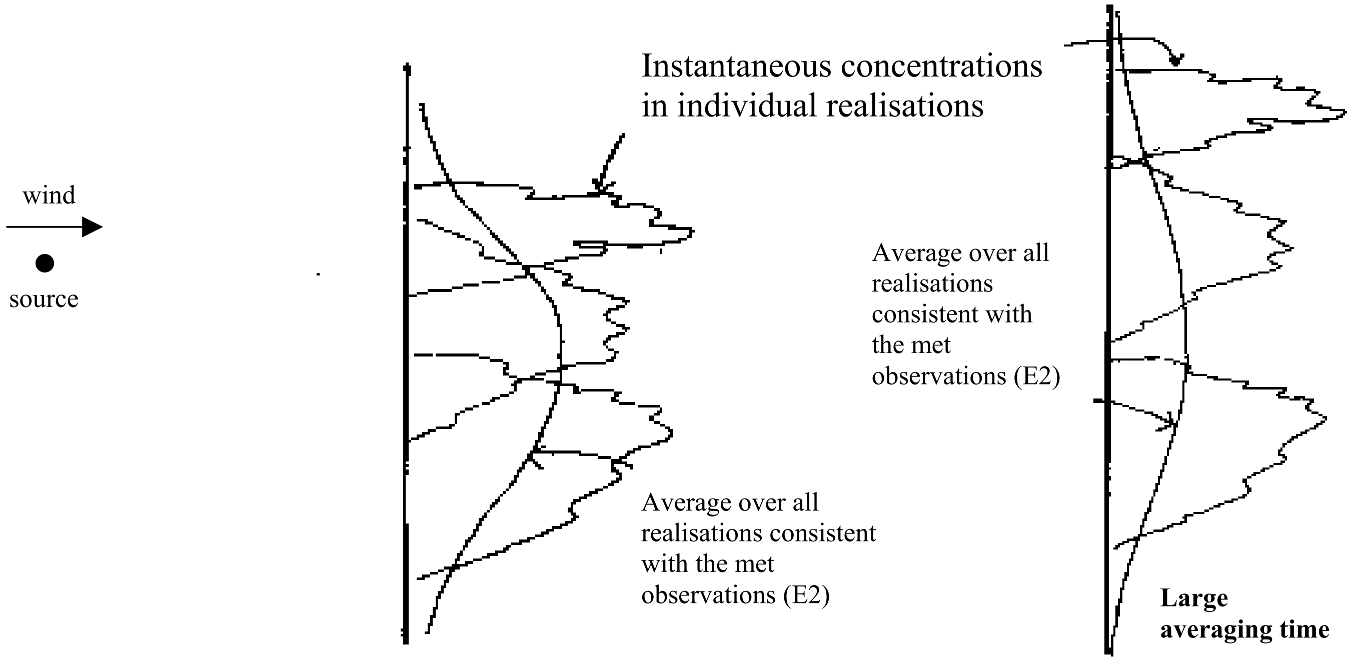

Large averaging time

Small averaging time

Figure 1b E2 averaging with a simple set of met data

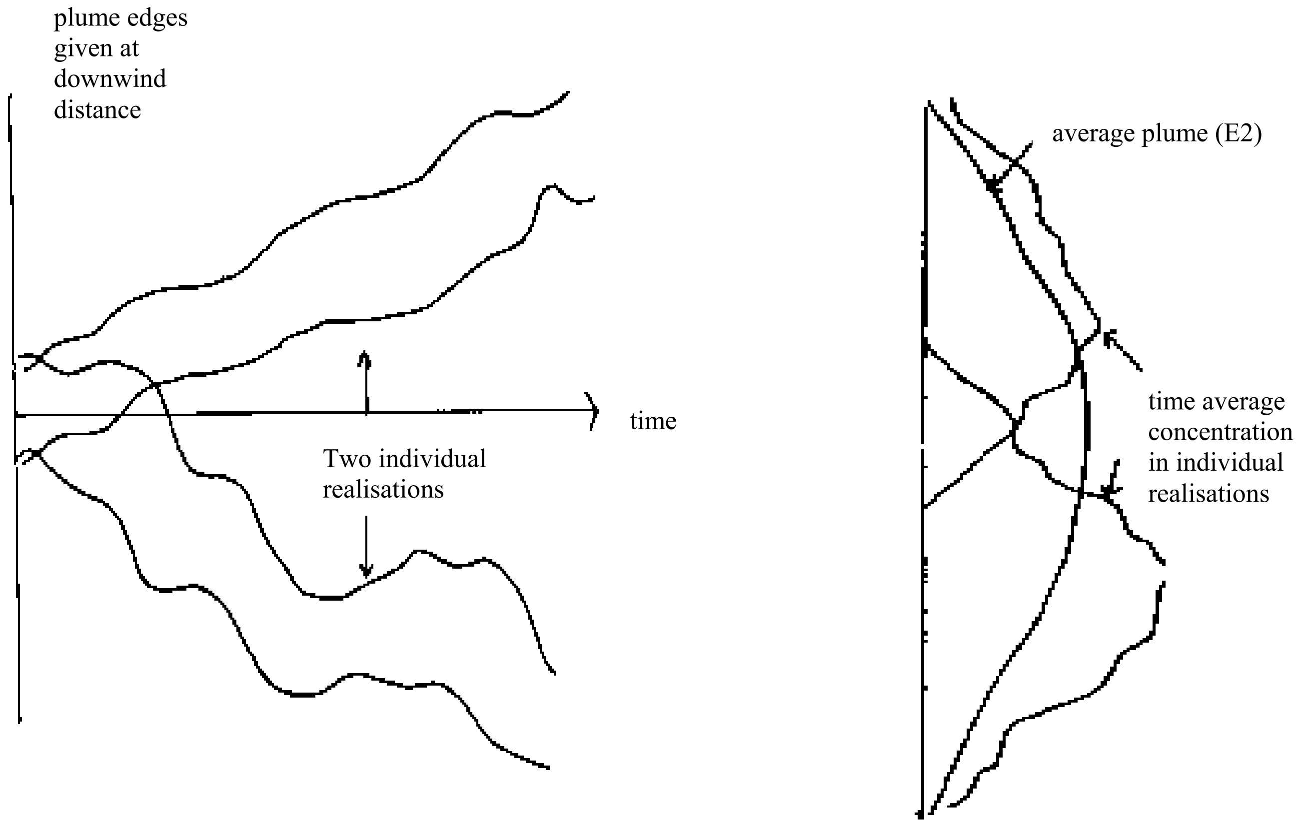

Figure 2a Shows the case where the met (in particular the wind direction) is appropriate to the start of the release period

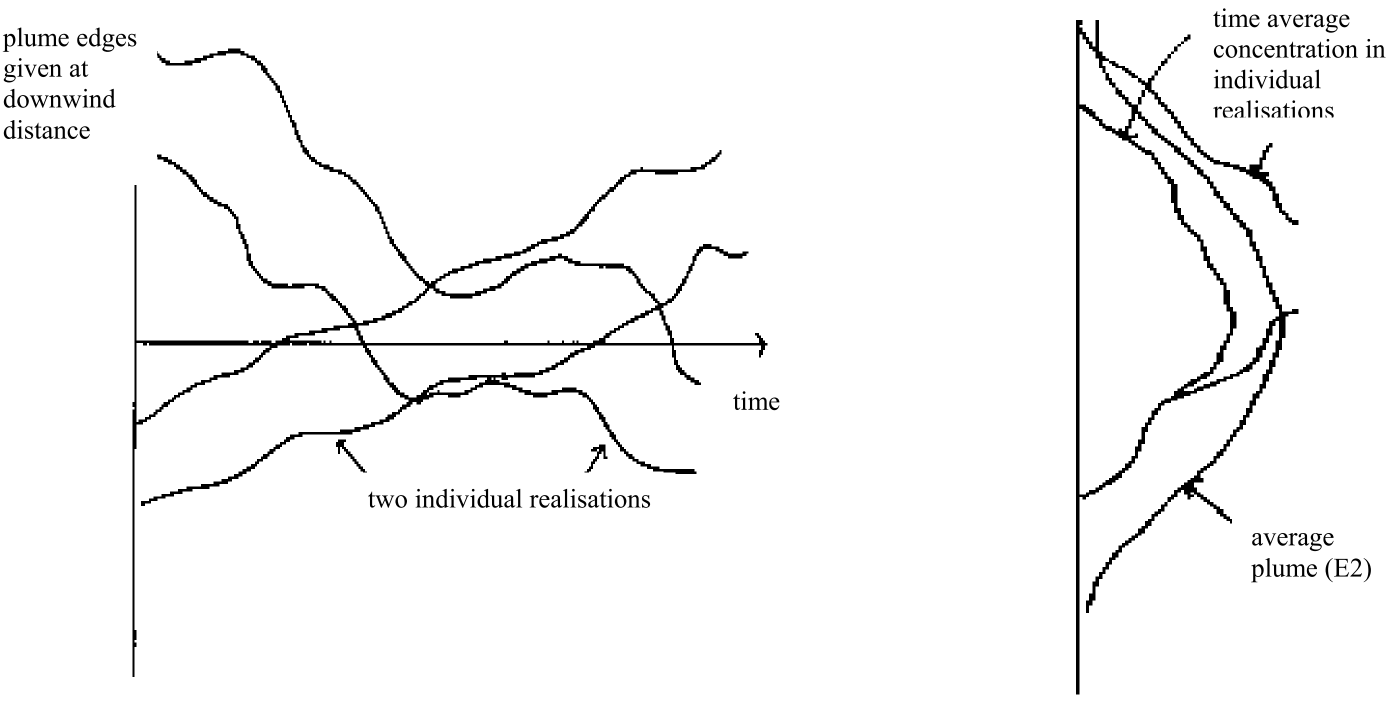

Figure 2b Shows the case where the met (in particular the wind direction) is appropriate to the middle of the release period

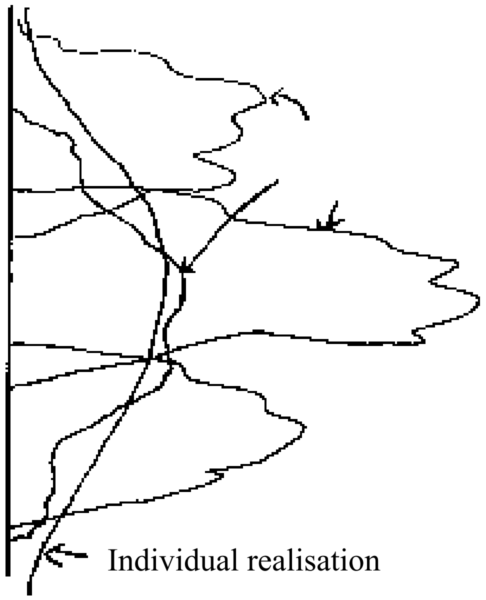

Figure 3a No centroid alignment

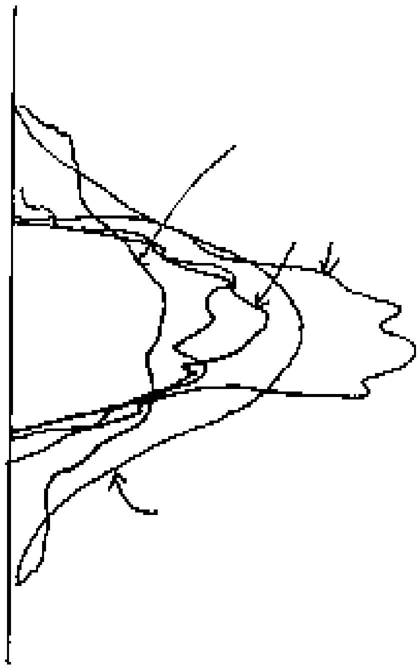

Figure 3b Centroid alignment

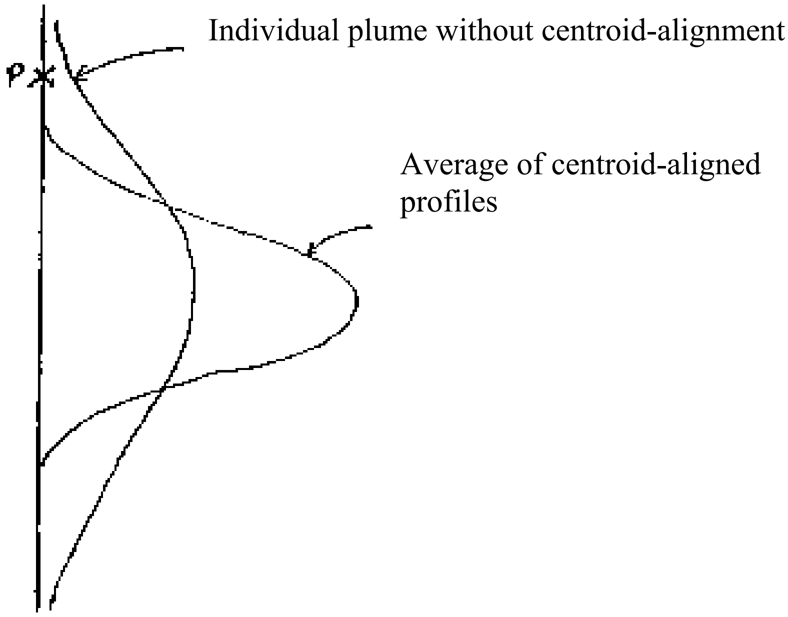

Figure 3c Comparison of average profiles. P indicates a point where there is a risk of encountering the plume which might be neglected by users if the centroid aligning is adopted.

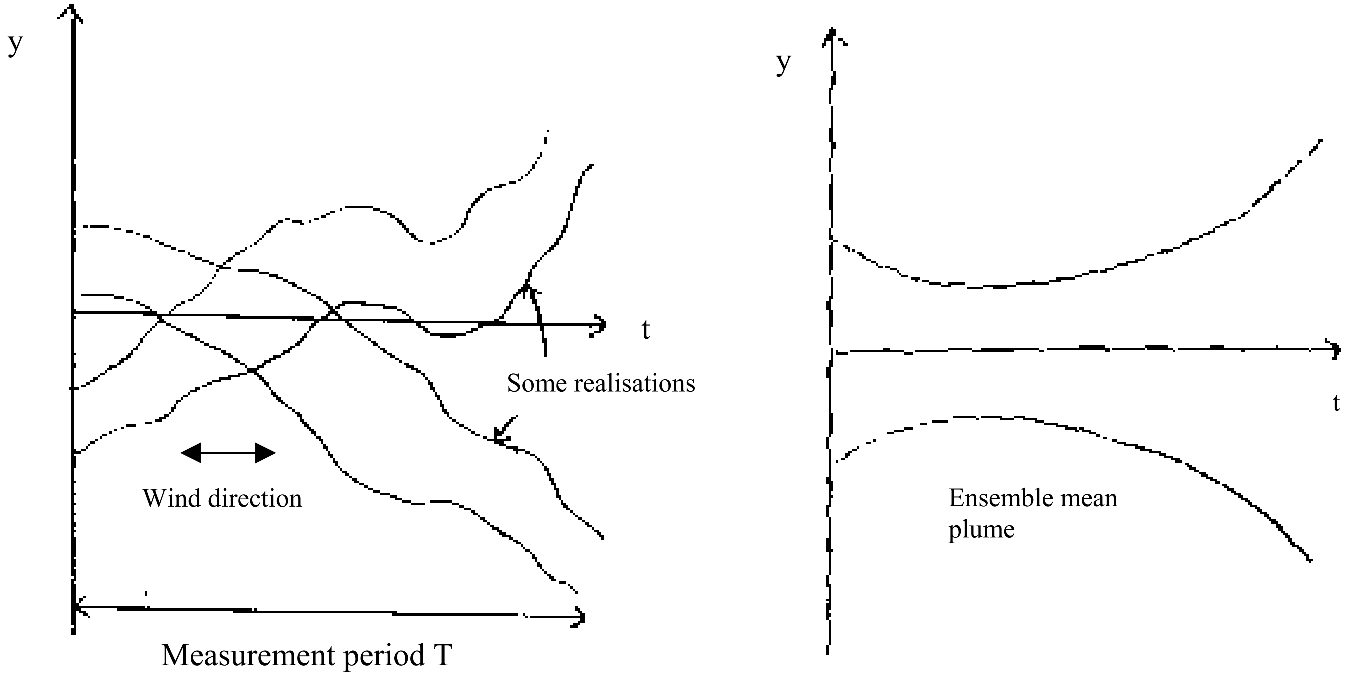

Figure 4a Wind direction measured over a short period

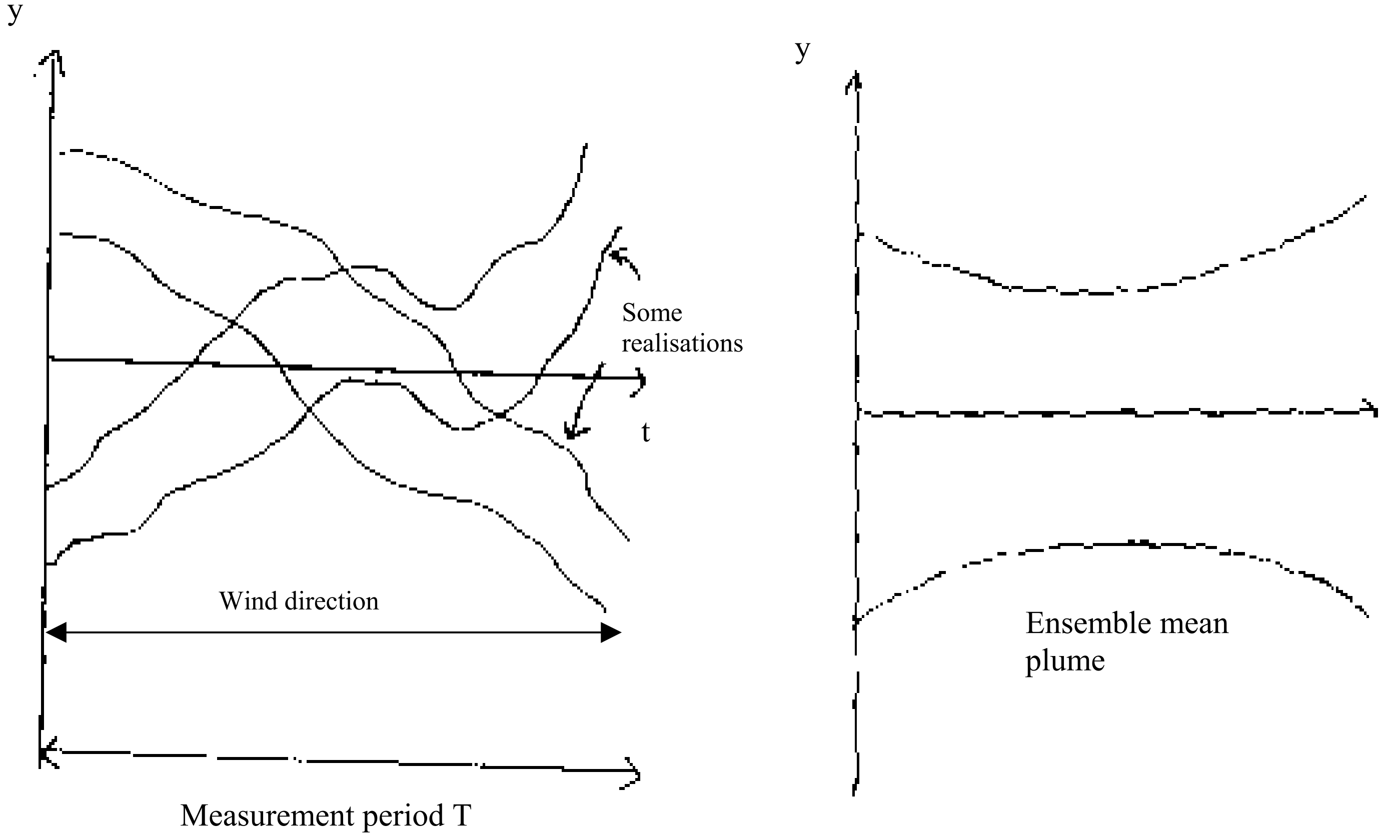

Figure 4b Wind direction measured over period T

Figure 4 Many measurements of length t_{av} over a period T (t_{av}«T) for a continuous source. y = 0 is the direction of the wind measurement.