Available online at www.sciencedirect.com
SCIENCE DIRECT®

Aerosol Science 36 (2005) 846–865

Journal of
Aerosol Science

www.elsevier.com/locate/jaerosci

A modified hygroscopic tandem DMA and a data retrieval method based on optimal estimation

M.J. Cubison*, H. Coe, M. Gysel

School of Earth, Atmospheric and Environmental Science, University of Manchester, Sackville Street Building, PO Box 88, Manchester, M60 1QD, UK

Received 24 May 2004; received in revised form 15 November 2004; accepted 17 November 2004

Abstract

We present a new design of hygroscopic tandem differential mobility analyser (HTDMA) capable of rapidly switching between measurement at constant humidity and the determination of humidograms. Both technical aspects of the instrumental setup and the software control are discussed.

The optimal estimation method (OEM) developed by Rodgers (Rev. Geophys. 14 (1976), J. Geophys. Res. 95 (1990), Inverse Methods for Atmospheric Sounding, World Scientific, Singapore 2000) for solving atmospheric data inversion problems has been successfully developed to retrieve measurements of hygroscopic particle growth using the tandem DMA. The technique makes no assumptions about the shape of the hygroscopic growth distribution and is able to determine the resolution and shape of the retrieved distribution. The technique has been shown to be robust throughout extensive tests and a thorough error analysis has been performed. The largest source of error arises from the under-sampling of the measurements, the so-called 'smoothing error'. However, the uncertainties in the measurements and error in the forward model are also significant. The technique allows hygroscopic growth measurements to be retrieved under a variety of atmospheric conditions and over an extended time period in a reliable and consistent manner.

A comparison is made between the OEM routine and the well-documented TDMAfit method of data analysis, showing the limitations and subjectiveness of TDMAfit in certain situations.

* Corresponding author. Tel.: +44 161 236 3311 X2490; fax: +44 161 200 3951.
E-mail addresses: m.cubison@postgrad.manchester.ac.uk (M.J. Cubison), hugh.coe@manchester.ac.uk (H. Coe), m.gysel@manchester.ac.uk (M. Gysel).

doi:10.1016/j.jaerosci.2004.11.009

The OEM data retrieval process is a transferable routine that could also be used in other applications, such as the volatility tandem DMA instrument.

Keywords: Hygroscopicity; DMA; Instrument; Inversion

1. Introduction

The hygroscopicity of Aitken and accumulation mode particles can be investigated as a function of size using the hygroscopic tandem differential mobility analyser (HTDMA) (Rader & McMurry, 1986). In the HTDMA, a quasi-monodisperse distribution of dry particles (<10% RH) is first selected from a polydisperse distribution using a differential mobility analyser (DMA). The monodisperse distribution of particles is then exposed to a controlled humidity environment before being passed into a second DMA. If the particles take up water then they will grow in size in the humid environment—the second DMA is stepped over an appropriate size range and the output particles are counted using a condensation particle counter (CPC). A plot of particle diameter from the second DMA versus CPC particle count will then give a distribution which determines the size to which the original monodisperse population has grown. In addition, through careful control of the humidity to which the particles are exposed between and in the DMA's, it is possible to construct a relationship between the growth and the relative humidity (RH) for a given monodisperse input size, known as a humidogram. The ratio of the wet size of an aerosol particle to its dry size is known as the growth factor, and is usually defined for a given humidity.

Recent developments in HTDMA design have seen some marked improvements on earlier instruments, including the low-temperature design developed by Weingartner, Gysel, and Baltensperger (2002) and the principles of stable operation reported by Maßling, Wiedensohler, and Busch (1999). Cocker, Flagan, and Seinfeld (2001) study the growth of secondary organic aerosol in an artificial chamber using a system which is regulated to a temperature identical to that of the chamber. These studies have demonstrated both the importance of maintaining steady RH in the HTDMA and reported designs capable of achieving this. However, rapid alteration of the measurement humidity has not yet been reported in the literature.

We present a new design of HTDMA capable of rapidly switching between measurement at constant humidity and the determination of humidograms, including both the deliquescence and efflourescence curves of the ambient aerosol.

Many previous HTDMA measurements have been analysed by a gaussian fitting method, pioneered by Stolzenberg and McMurry, who developed the TDMAfit program (TDMAfit users manual, Particle Technology Laboratory, Department of Mechanical Engineering, University of Minnesota, USA, ptl publications no. 653, 1988.) This assumes the transfer function of the instrument is of a gaussian shape, and then fits one or more gaussian peaks to the raw data set to determine the peak positions and relative concentrations of the modes (e.g. Zhou, Swietlicki, Hansson, & Artaxo, 2002; Swietlicki et al., 1999; Higgins, Jung, Kittelson, & Roberts, 2002). Other studies have solved a set of linear equations to invert the measurement data (e.g. Vakeva, Kulmala, Stratmann, & Hameri, 2001). However, these methods are not always satisfactory, as it is not known if the separation of the modes is less than the resolution of the data set, and it does not give an opportunity for a thorough error analysis. When solving linear equations the problem is often not well-posed and the solution space can contain many possible outcomes.

Here we describe a new retrieval method that inverts growth factor measurements, but the relative proportions of particles in a range of growth factor ‘bins’ are retrieved, along with the associated errors. In addition, the retrieval process gives a measure of the resolution of the retrieved data, and therefore the minimum mode separation that can be resolved.

2. Experimental setup

The HTDMA used in this work, based on the design by Weingartner et al. (2002) consists of a bipolar neutraliser, a fixed voltage long column DMA that is used to select a narrow size range of ambient aerosol, a humidification system, a second DMA and a CPC (CPC model 3760, TSI Inc.) used to determine the growth of the particles selected by the first DMA. The DMAs used were of the Vienna design (see Winkelmayr, Reischl, Lindner, & Berner, 1991), with a central column length of 0.28 m and sheath flow rate 6 litres per minute (lpm).

A diagram showing the layout of the HTDMA can be seen in Fig. 1.

2.1. Size calibration

Particle mobility size measurements made using the DMAs were referenced to a primary standard through calibration using NIST-traceable polystyrene latex spheres (PSLs). These provided size calibration for particle diameters greater than 90 nm. PSL’s smaller than this threshold failed to successfully

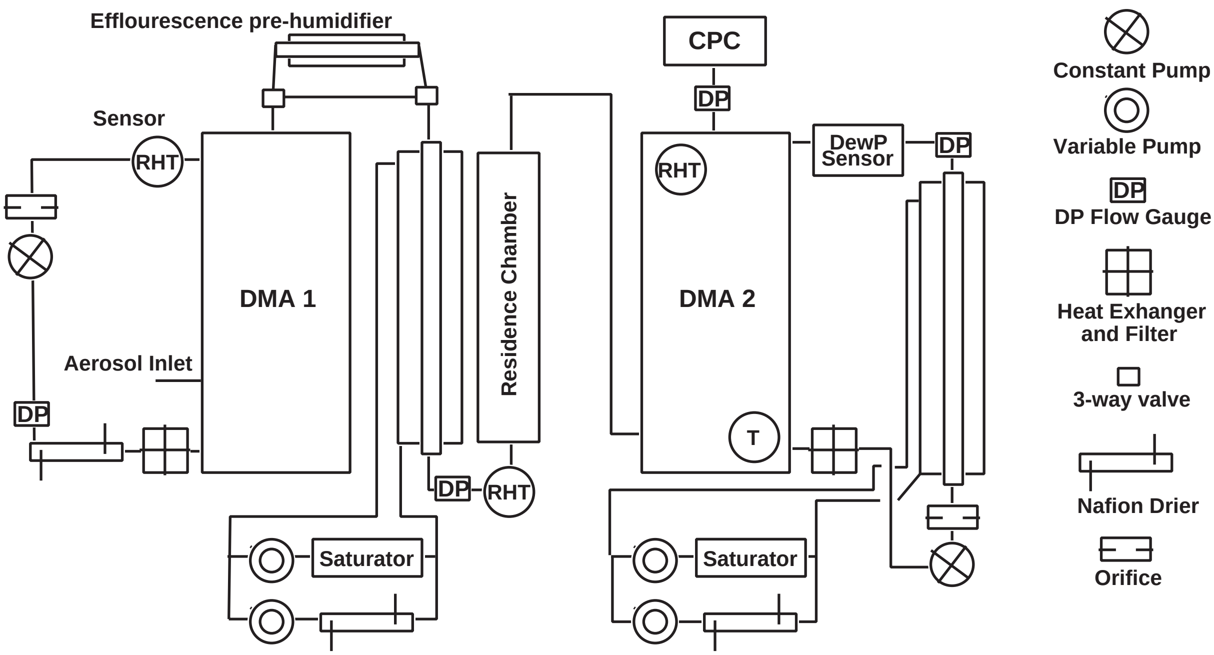

Fig. 1. Schematic diagram showing the layout of the HTDMA described in this work.

nebulise in the atomiser, thus it was necessary to use another method of calibration for the smaller diameters. This was achieved by running the second DMA at a low value of RH (less than 15%), and comparing the offset of the two DMAs at a range of sizes. Together with the PSL calibration this defines the percentage offset in the two DMA's over the entire particle size range measured.

2.2. Re-circulating sheath flow

A re-circulating sheath flow system was used in both DMAs in order to facilitate accurate control of both flow rate and RH. Total particle HEPA filters were used in both DMAs to ensure clean, particle-free sheath air was injected at the bottom of the DMAs. Differential pressure gauges (Omega 0.25/1″H20) gave an instantaneous measure of flow rate, and were calibrated using a Gilian primary flow calibrator. In DMA1 the re-circulating sheath flow circuit also contained a Nafion dryer (PermaPure, model PD100T 12″long) capable of maintaining $<10\%$ even when the ambient humidity is high. In the sheath circuit for DMA2 the sheath air drier was replaced by a humidifier, details of which shall be discussed in Section 2.3. In both DMAs, the sheath flow was heat exchanged to the temperature of the DMAs immediately before injection into the DMA. The pumps (KNF 026, Teflon Head) were thermally insulated from the DMAs, and the flow rates controlled by use of a critical orifice.

2.3. Humidification

Humidification of the aerosol and sheath lines was achieved using separate, independent humidifying circuits. This allows rapid adjustment of RH, and also determination of the hysteresis curve of the growth behaviour of particle hydration and dehydration. In both cases, the flow was passed through a humidifier consisting of 0.5 m length of Gore-Tex© tubing, around which a counter-flow of humid air was passed. The RH of this purge flow was determined by two variable speed, computer controlled pumps, which mixed two lines of saturated and dry air to achieve a purge airstream of given humidity. The control software monitored the aerosol and sheath RH via a Rotronic capacitative probe and the dew point sensor respectively, and adjusted the ratio of the variable speed pumps to maintain these values to a given setpoint. The sheath and aerosol humidifiers were served by separate purge mixing circuits, allowing independent control of sheath and aerosol RH.

In the case of the aerosol humidifier, the inner flow rate was just 0.6 lpm, so a change in purge RH rapidly altered the aerosol RH also.

In the sheath humidifier, the flow rate was 6 lpm, so despite a higher purge flow, the response was much slower. However, in the re-circulating sheath circuit, the humidifier only needs to make small changes to the sheath flow water vapour content in order to be extremely effective, as the airstream is continuously pumped around and back through the humidifier for further RH adjustment.

In addition, a third humidifier was incorporated into the system. This was named the efflourescence humidifier (EH), and contained a purge flow of water rather than air. The aerosol flow was passed through this just before entering the aerosol humidifier when measuring the efflourescence curve of the aerosol. With a flow rate of 0.6 lpm and initial humidity of under 10%, the output from the EH was in excess of 90% RH. To measure the efflourescence (dehydration) curve of the aerosol, it was first passed through the EH, and then dried to a particular value using the aerosol humidifier, before being passed into the second DMA, which was also maintained at the same value of RH, for determination of growth factor.

2.4. Temperature control

Variations in temperature are of particular concern when making HTDMA measurements, as a shift of only one degree can bring a change in RH of up to 6% at high humidities. Hence it is critical to maintain the DMAs at a constant temperature to avoid fluctuations in RH. To try and achieve this the DMAs were surrounded by water piping, around which a constant water flow of 2 lpm was pumped from a central 10 l tank. Both DMAs were then lagged using glass fibre insulation. The aerosol and sheath lines were heat exchanged to the tank temperature prior to entering DMA2, and all heat sources, such as pumps and power supplies, were kept in a separate, insulated unit.

All the temperature and RH readings were continuously recorded allowing corrections and diagnostics to be performed at a later date. Current work is focusing on eliminating temperature fluctuations in the instrument by submersing both DMA's in two separate water tanks, with independent temperature control provided by peltier units.

2.5. RH/T measurement

Primary RH measurement was obtained using a General Eastern Hygro M4/D-2 dew point sensor located on the sheath excess outlet of the second DMA, this was capable of measuring RH to within 1.5%. Several capacitative probes (Rotronic, SC05) were used in-line throughout the instrument to monitor temperature and RH, and were calibrated against the dew point sensor. These probes were mounted within Swagelok© cross pieces, and in-line losses were shown to be negligible through measurements taken with a Scanning Mobility Particle Sizer (TSI, Model 3080).

Calibration of the probes was achieved by running a humid airstream through both the dew point sensor and the probes, whilst keeping both in a thermostatically stable environment. The process was repeated at a range of humidities from 10 to 95% RH and a mathematical fit applied to the resulting difference curve. The probes were found to under-predict the RH by as much as 6% at high humidities of 90% RH and above. In addition, the difference between the probes and the primary dew point sensor standard tended to increase with time. It is therefore necessary to carry out the calibration process immediately before and after field and laboratory campaigns.

3. Software control

Control software for the HTDMA was encoded using National Instruments LabVIEW version 6, and contained three main elements, a program to monitor all the output variables from the HTDMA, some feedback routines to control the variable speed pumps depending on these outputs, plus a logging program to control the power supplies and retrieve data from the CPC.

The analogue output variables from the HTDMA were read at 40 Hz, with the data returned from the monitoring program giving 1 s averages. Data was stored as 10 s averages, which is an appropriate resolution for post-experimental diagnostics, as the response of the instrument to RH changes is on this time scale. In addition, during scanning, the dew point sensor was continuously accessed via the serial port, and returned 6 s averaged data.

The logging program ran independently of the control and monitoring software, and was set up to using a stepping-voltage system to measure the growth factor peaks in DMA2. In some previous HTDMA systems the scanning-voltage approach has been taken in the second DMA, but the stepping method gives better counting statistics and therefore reduced errors. However, the stepping method offers considerably less temporal resolution than the scanning method. In this system, measurements were made at 7 different dry sizes every hour, each with approximately 30 different sizes selected in the second DMA. In times of low ambient particle concentration the counting statistics were improved by waiting longer between voltage steps and therefore averaging over a longer time period. Fewer data points were taken in this situation in order to maintain the one-hour total scan time. Growth factor measurements were made from 0.85 to 2.45 at all the different dry sizes.

Humidograms were determined using a routine included in the constant humidity logging program. This routine used the same logging software to make the measurements at individual humidities and invoked the control software to rapidly alter the RH between measurements. Both deliquescence and efflourescence measurements were made at each humidity, usually for 2 different dry aerosol sizes. Electronically-actuated valves switched between the 'conventional' deliquescence, or hydration, measurements and those of the efflourescence, or dehydration, curve.

Control of RH in the aerosol line and in DMA2 was achieved through a feedback loop taking data from either a Rotronic sensor (aerosol line) or the dew point sensor (DMA2). The variable speed pumps controlling the RH of the purge flow in the humidifiers were then altered accordingly using a specially constructed control loop taking into account the long response time of the instrument. Using this system, the aerosol RH could be controlled to within 0.2% RH of the setpoint, and the sheath RH (in DMA2) to within 1%, which is within the error of the RH measurement by the dew point sensor.

The total time to make a complete determination of the hysteresis curve at two different sizes, and return to 90% RH for continuation of constant humidity scanning, was 2 h. An example of the capability of the instrument to change RH is demonstrated in Fig. 2. When the setpoint RH is achieved, the instrument

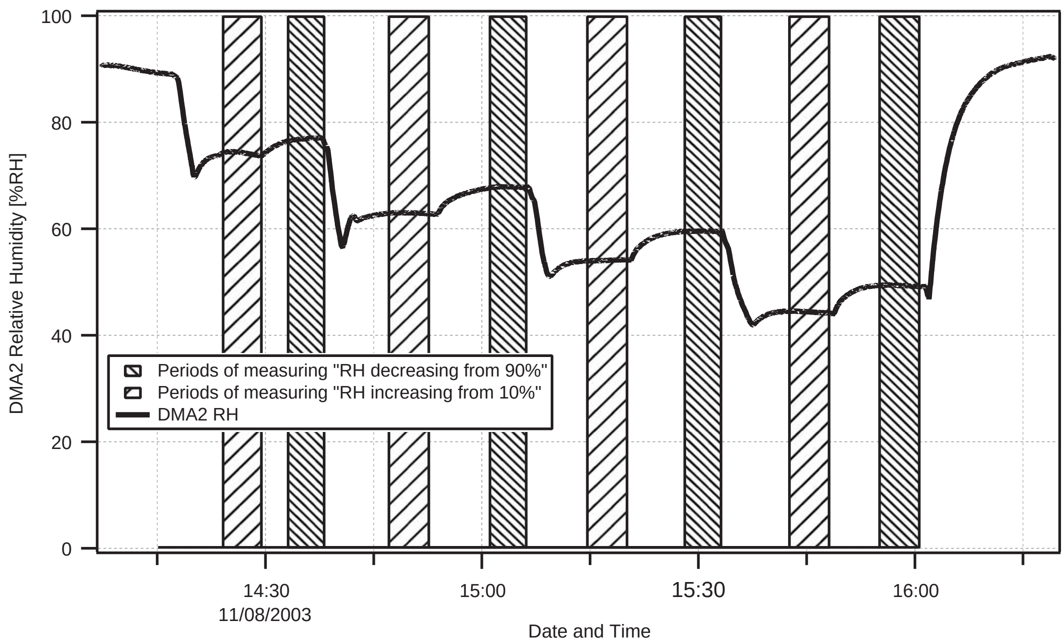

Fig. 2. Example of sheath RH response during a humidogram taken on the TORCH experimental campaign.

takes a reading, as indicated by the shaded areas on the diagram. The monitoring software was capable of measuring the RH to the instrumental accuracy of 1.5% RH. The control software could maintain a setpoint RH within this range as indicated above. For a change from 90% RH to 10% RH, the aerosol control circuit can achieve dry aerosol conditions within 1 min, and the sheath circuit within approximately 2 min.

4. Data retrieval and analysis

In the HTDMA, the final aerosol distribution after the second DMA represents the range of growth factors of the dry aerosol leaving DMA1 convolved with the inherent broadening of the instrument due to the transfer function. To separate these two effects and assess the true growth of the aerosol it is necessary to perform some post-measurement processing of the experimental data.

Many previous HTDMA studies have used data-inversion routines based on the principles of the TDMAfit program developed by Stolzenburg and McMurry. The measured growth factors are assumed to fall into a few, well defined separated modes, which are log-normally distributed. They are described by an arithmetic mean given by the growth factor value and the standard deviation of the mode, multiplied with the value of a second calculated parameter, the growth dispersion factor. In many cases, these parameters were determined by fitting 1–3 normal distributions of diameter growth factors to the measurement data.

However, the gaussian fitting method does not give a detailed error analysis on the data set, nor does it give a measure of the resolution of the measurements. It is thus possible to apply a bi-modal gaussian fit to a measurement distribution where the splitting of the modes is finer than the resolution of the measurements. In addition, when conducting hygroscopic closure studies, it is desirable to combine the particle size distribution, composition and hygroscopic growth factor distribution to achieve closure on the data set. However, the gaussian fitting method does not return a growth factor distribution, but merely the positions and widths of the growth factor modes.

We present a novel method of inverting HTDMA measurement data, using the optimal estimation method (OEM) for solving atmospheric data inversion problems developed by Rodgers (see Rodgers, 1976, 1990, 2000), hereafter referred to as the OEM. The approach followed here retrieves a distribution in growth factor space from the measurement distribution using an empirically determined transfer function of the instrument. It allows a full error analysis and gives a measure of the optimum resolution of the inversion.

The OEM offers several important advantages over other inversion methods. Firstly, inversion problems of this type are frequently very poorly constrained and the solution space offers several possible outcomes. OEM uses an a priori distribution to constrain this space. Secondly, it makes an explicit error analysis, allowing the individual contributions to the overall error to be assessed. Lastly, it involves no large iterations and is therefore very computationally efficient.

The HTDMA provides us with a series of particle concentration measurements at different wet diameters for a given dry diameter. However, a plot of growth factor, without the effect of the transfer function of the DMAs, including diffusional and atmospheric broadening, cannot be easily calculated. Despite this, given a true growth factor distribution we should be able to calculate the measurement distribution returned using a forward model:

$$\mathbf{y} = \mathrm{F}(\mathbf{x}, \mathbf{b}) + \varepsilon_{\mathbf{y}}, \tag{1}$$

where $\mathbf{y}$ is a vector of dimension m representing the HTDMA measurements, $\mathbf{x}$ the true growth factor distribution, given for a finite number, n, levels, and $\varepsilon_{\mathbf{y}}$ the measurement errors corresponding to the appropriate elements of $\mathbf{y}$. F is the forward model containing several controlling parameters, $\mathbf{b}$. In the real atmosphere $\mathbf{x}$ is continuous, but the inversion procedure takes it to be discrete. The separation of these discrete elements can be varied, but the information content is limited by the smoothing function of the inversion procedure, which shall be discussed later.

The inverse procedure of Eq. (1) would return the true growth factor distribution from HTDMA measurements, which is the process we wish to carry out. However, the problem is under-constrained, therefore there is no unique solution for any given set of measurements. Measurement errors, given by $\varepsilon_{\mathbf{y}}$, will further reduce the information content available in the retrieved distribution, and contribute to what is known as the *null space*. This is the hidden information that cannot come from the measurements owing to their finite resolution in growth factor space. Any information about the null space must therefore come from a priori knowledge of the true growth factor distribution.

5. Mathematics of the retrieval process

Eq. (1) gives a relationship between the HTDMA measurements, $\mathbf{y}$, and the true growth factor distribution of the atmosphere, $\mathbf{x}$, on which they depend. In the process of data retrieval we wish to calculate the inverse of this relationship along with its associated errors.

For the purpose of such an error analysis we can consider linear perturbations of the forward model about a reference state $(\mathbf{x_0}, \mathbf{b_0})$, which may be the true growth factor distribution, $\mathbf{x}$, the retrieved distribution, defined as $\hat{\mathbf{x}}$, or any other arbitrary state as required.

$$\mathbf{y} = \mathrm{F}(\mathbf{x_0}, \mathbf{b_0}) + \frac{\partial \mathrm{F}}{\partial \mathbf{x}}(\mathbf{x} - \mathbf{x_0}) + \frac{\partial \mathbf{F}}{\partial \mathbf{b}}(\mathbf{b} - \mathbf{b_0}) + \varepsilon_{\mathbf{y}}. \tag{2}$$

However in the HTDMA the problem posed is linear, thus the weighting function matrix, defined as $\mathbf{K}{=}\partial\mathbf{F}/\partial\mathbf{x}$, is independent of x. The matrix $\mathbf{K}$ also represents the forward model of the instrument, where each row corresponds to a different measurement and each element along a row represents the sensitivity of each measurement to the different distribution elements.

In addition, $\mathbf{K_b}$ is defined to be the sensitivity of the measurement to the forward model parameters, $\partial\mathbf{F}/\partial\mathbf{b}$. These derivatives may be evaluated numerically by perturbing the forward model, and is the approach used here. This is discussed further in Section 8.

To obtain a formula for the true distribution, OEM combines the measurements with an a priori distribution in such a way that their contributions are weighted according to their variances. To take the simplest case, consider two measurements x_1, x_2 of a variable x, with corresponding variances S_1 and S_2. The best estimate of x is then:

$$x_{\mathrm{av}} = \frac{S_1^{-1}x_1 + S_2^{-1}x_2}{S_1^{-1} + S_2^{-1}}. \tag{3}$$

In the OEM, the matrix analogue of this equation is used to combine the measurement distribution, $\mathbf{y}$ and its associated covariance matrix, $\mathbf{S}_\varepsilon$, with the a priori. The a priori can be considered simply as a virtual measurement, $\mathbf{x_a}$, with an associated covariance matrix $\mathbf{S_a}$. From Eq. (2), in the linear case we can choose any reference state $\mathbf{x_0}$. By choosing $\mathbf{x_0} = 0$ and assuming there are no measurement errors or systematic errors from the forward model, $\mathbf{y} = \mathbf{Kx}$. Therefore, $\mathbf{x} = \mathbf{K}^{-1}\mathbf{y}$ gives an estimate of the true distribution, $\mathbf{x}$. The inverse covariance of the measurements is then given by $\mathbf{K}^{\mathbf{T}}\mathbf{S}_{\varepsilon}^{-1}\mathbf{K}$. Thus, taking the matrix analogue of Eq. (3):

$$\hat{\mathbf{x}} = \frac{\mathbf{S}_{\mathbf{a}}^{-1}\mathbf{x}_{\mathbf{a}} + \mathbf{K}^{\mathbf{T}}\mathbf{S}_{\varepsilon}^{-1}\mathbf{y}}{\mathbf{S}_{\mathbf{a}}^{-1} + \mathbf{K}^{\mathbf{T}}\mathbf{S}_{\varepsilon}^{-1}\mathbf{K}} \equiv (\mathbf{S}_{\mathbf{a}}^{-1} + \mathbf{K}^{\mathbf{T}}\mathbf{S}_{\varepsilon}^{-1}\mathbf{K})^{-1}(\mathbf{S}_{\mathbf{a}}^{-1}\mathbf{x}_{\mathbf{a}} + \mathbf{K}^{\mathbf{T}}\mathbf{S}_{\varepsilon}^{-1}\mathbf{y}). \quad (4)$$

In practice, the retrieved distribution, $\hat{\mathbf{x}}$, will not be the same as the true distribution, $\mathbf{x}$, the exact values of which are not known. However, OEM provides a powerful tool to make the best estimate of this true distribution possible with the available measurements, whilst giving an explicit evaluation of the errors involved in the process.

6. The forward model

In Eq. (1) we considered a forward model F(**x**,**b**), where **b** is a vector of model parameters that are not perfectly known to the observer, and hence a possible source of random or systematic differences between calculated and measured values of **y**. It is important that the forward model should contain all the physics of the experiment, to an accuracy greatly exceeding that of the measurement. If this is not the case, then the accuracy of the model should be quantifiable, otherwise the retrieved distribution will not be accurate and the measurements of no real value.

In the case of the HTDMA, the forward model is calculated by considering the effect of the two DMAs and the humidifier on the input aerosol distribution.

In order to ascertain the exact transfer function of each DMA, mathematical fits were made to calibration curves made using latex spheres (PSLs) of a known size-distribution. It is then possible to eliminate the width arising from the PSLs in the measurement data, leaving only the transfer function of the DMA, as shown by the dotted line in Fig. 3. This was done for each DMA and the overall transfer function of the instrument at a diameter d_{tr} and growth factor g_{tr} is then a combination of the first DMA transfer function

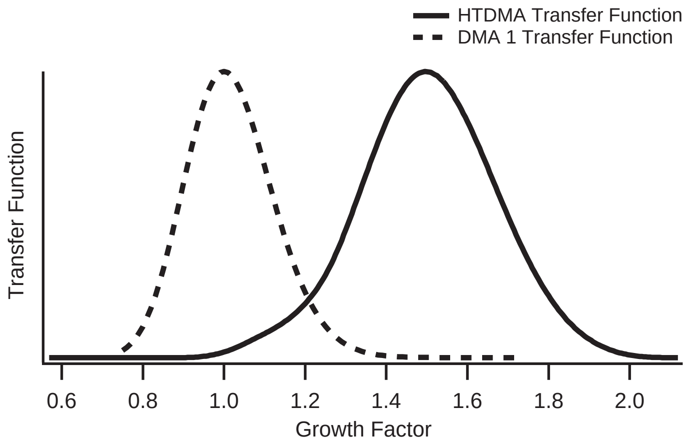

Fig. 3. The combination of two DMA transfer functions gives the overall HTDMA transfer function (forward model width). Here the growth factor of the particles is 1.5 in diameter space.

at d_{tr}, the transfer function of the humidifier for particles of growth g_{tr} and the second DMA transfer function at diameter $d_{tr}g_{tr}$. This is shown by the bold line in Fig. 3. We can calculate $F(g_0, g)$, the output width of the forward model over the growth factor range g arising from a delta function at position g_0, for g_0 at each retrieved point, which gives the weighting function, $\mathbf{K}$.

Using this method, the exact mathematical curve describing the transfer function may change with each calibration of the instrument. Factors such as the sheath and aerosol flows, the particle diameter measured and contamination of the DMA's can affect the transfer function of the DMA's. Any such variability in different calibration curves gives a measure of the error in the forward model.

7. The a priori distribution

In this work, the a priori distribution, as shown in Fig. 4, is composed of a range of previous atmospheric HTDMA measurements available in the literature:

- McMurry and Stolzenburg (1989)—Los Angeles, USA.
- Zhang, McMurry, Hering, and Casuccio (1993)—Los Angeles and Grand Canyon, USA.
- Covert and Heintzenberg (1993)—Svalbard, Norway.
- Svenningsson and Hansson (1992)—Po Valley, Italy.
- Svenningsson and Hansson (1994)—Kleiner-Feldburg, Frankfurt, Germany.
- Svenningsson et al. (1997)—Great Dun Fell, UK.
- Pitchford and McMurry (1994)—Grand Canyon, USA.
- Fuzzi et al. (1998)—Po Valley, Italy. CHEMDROP campaign.
- Maßling, Wiedensohler, and Voutilainen (2003)—ACE-Asia campaign.

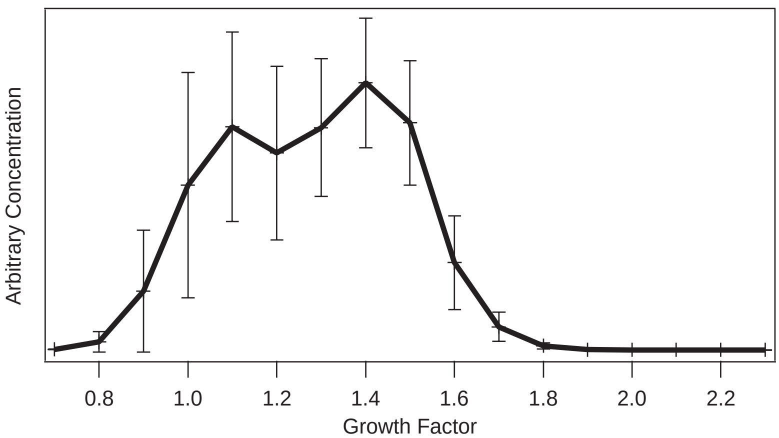

Fig. 4. The synthetic a priori, $\mathbf{x_a}$, and associated errors used in the example retrieval. It is important to choose a distribution that reflects the measurements being taken—for example, the distribution shown would be inappropriate for use in retrieving measurements taken in the marine boundary layer, as it contains no information from studies of this type. As such the observed high growth mode arising from sea salt is not contained in the a priori distribution.

The various sets of data have been averaged across their growth factor bin to give an average a priori distribution, $\mathbf{x_a}$. This gives a climatological average distribution and the associated variances from a wide range of field data. It is important to select a a priori distribution that reflects the problem and therefore offers some information to the retrieval process. In this work the urban and background studies chosen to construct the a priori reflect the nature of the site at which measurements were taken that are used here to test the OEM.

Where measurement data are sufficient, the final retrieved distribution will only depend weakly on the a priori distribution, as shown in Section 8. However, OEM uses the errors in the a priori distribution as well as the distribution itself to constrain the retrieval process. These errors are represented in Eq. (4) by the variances of the column vector $\mathbf{x_a}$ along the diagonal of matrix $\mathbf{S_a}$, whose remaining elements are zero if the covariances between the individual a priori elements are negligible. Hence large errors in the a priori distribution will lead to the first term in the numerator of Eq. (4) being small in comparison to the second term arising from the measurements. In this case the a priori distribution lends very little to the retrieval. Conversely, large errors in the measurement data and/or very small errors in the a priori distribution will mean that the measurements add no further information to that already available in the literature.

Occasionally, the retrieval may be improved by a second iteration of the OEM using the result of the first iteration as the a priori for the second. This is usually done in the case where the climatological a priori is far from the ideal solution, which emphasises the importance of selecting an appropriate ensemble of distributions from which to create the a priori. The errors from the first iteration must still be used in the second iteration, as these cannot be reduced simply by iterating the retrieval process.

8. Errors in the retrieval

Rodgers (2000) describes two tools which are particularly useful when considering the performance of a retrieval made using the OEM. These are the contribution function matrix, $\mathbf{G}$, and the averaging kernel matrix $\mathbf{A}$. The columns of the contribution function matrix describe the sensitivity of the retrieved distribution, $\hat{\mathbf{x}}$, to perturbations in a particular element of the measurement vector $\mathbf{y}$.

The contribution function matrix is defined as

$$\hat{\mathbf{x}} - \mathbf{x_a} = \mathbf{G}(\mathbf{y} - \mathbf{K}\mathbf{x_a}) \Rightarrow \mathbf{G} = \mathbf{S_a}\mathbf{K^T}(\mathbf{K}\mathbf{S_a}\mathbf{K^T} + \mathbf{S}_\varepsilon)^{-1}. \tag{5}$$

The averaging kernel matrix describes the contribution each point in the retrieved distribution makes to the true distribution at a given growth factor. It can be thought of as the response of the retrieval to a δ-function disturbance in the real distribution. A disturbance of this nature in the true distribution will be reflected in the retrieved distribution as a vector in a column of $\mathbf{A}$. In the ideal case, $\mathbf{A}$ would be a unit matrix, but in reality the columns and rows show peaked functions, the widths of which give a qualitative measure of the resolution of the retrieved distribution. $\mathbf{A}$ is defined as

$$\mathbf{A} = \mathbf{GK}, \tag{6}$$

where the rows of matrix $\mathbf{A}$ provide the averaging kernels.

The averaging kernel matrix is also a useful tool for detecting a misleading retrieval, if the shape is other than a simple peak. In particular, any attempts to retrieve a distribution with a growth factor resolution finer than the information content in the measurements will result in a retrieved distribution that displays oscillatory behaviour.

The total error in the retrieval is the difference between the retrieved concentration distribution, $\hat{\mathbf{x}}$, and the true distribution $\mathbf{x}$. Rodgers (1990) shows that this is the sum of three individual components:

$$\hat{\mathbf{x}} - \mathbf{x} = (\mathbf{A} - \mathbf{I})(\hat{\mathbf{x}} - \mathbf{x_a}) + \mathbf{G}\varepsilon_\mathbf{y} + \mathbf{G}\Delta_\mathbf{b}. \tag{7}$$

The first term on the right-hand side is the smoothing error. The averaging kernel described above implies that the retrieval process smoothes the true distribution and the difference between the two gives the smoothing error, which represents the parts of the real system over which the measurements provide no information. The smoothing error covariance matrix, $\mathbf{S_S}$, is given by

$$\mathbf{S_S} = (\mathbf{A} - \mathbf{I})\mathbf{S_a}(\mathbf{A} - \mathbf{I})^\mathbf{T}. \tag{8}$$

The second term on the right-hand side of Eq. (7) is the measurement error, representing the error in the retrieved distribution arising from the uncertainties in the growth factor measurements from the HTDMA. The measurement error covariance matrix $\mathbf{S}_\varepsilon$ is mapped into the retrieved distribution space using the contribution function matrix:

$$\mathbf{S_M} = \mathbf{G}\mathbf{S}_\varepsilon\mathbf{G}^\mathbf{T}. \tag{9}$$

The measurement error covariance matrix $\mathbf{S}_\varepsilon$ is formed by considering the contributions to the instrumental error for each growth factor measurement in the distribution. In the HTDMA there are three main contributions to this error. Firstly, the error arising from the counting statistics of the measurements, $\mathbf{S}_\varepsilon^{\text{counts}}$ can be represented according to Poisson statistics in the diagonal elements. The error in RH measurement, $\mathbf{S}_\varepsilon^{\text{RH}}$ becomes important if counting statistics are good. $\mathbf{S}_\varepsilon^{\text{RH}}$ is obtained by simulating the effect of a change in RH on the measured distribution, and it is not diagonal. Finally, there is an effect from the size calibration of the DMAs, which can be kept negligibly small through careful calibration procedures.

The last term on the right-hand side of Eq. (7) is the forward model error, representing the errors in the retrieved distribution that arise from uncertainties in the forward model. Rodgers (2000) formulates a forward model error covariance matrix that is mapped into the distribution space thus:

$$\mathbf{S_F} = \mathbf{G}\mathbf{K_b}\mathbf{S_b}\mathbf{K_b^T}\mathbf{G}^\mathbf{T}, \tag{10}$$

where $\mathbf{S_b}$ is the forward model parameter error covariance matrix and $\mathbf{K_b} = \partial\mathbf{y}/\partial\mathbf{b}$ represents the sensitivity of the growth factor measurements to the forward model parameters.

The matrix $\mathbf{K_b}$ is formed by perturbing each parameter in the forward model in turn and calculating the effect the perturbation has on the growth factor measurements. The matrices $\mathbf{K_b}$ and $\mathbf{G}$ are then used to map the forward model parameter errors into distribution space. The forward model in this case is a combination of the transfer functions of the two DMAs and the humidifier, as described in Section 6, and is sensitive to the position and widths of the mathematical descriptions of the DMA transfer functions. These parameters are used to formulate $\mathbf{K_b}$ and the squares of their uncertainties form the diagonal elements of $\mathbf{S_b}$. It is assumed the off-diagonal elements of $\mathbf{S_b}$ are zero, i.e. that the parameters are independent.

Fig. 5 shows the sensitivity of the growth factor measurements to these three parameters for the synthetic measurement distribution shown in Fig. 6. The largest error arises from the uncertainty in the position of the forward model peak, although the width and amplitude are also significant. To minimise these errors a number of calibration scans should be performed, to give a statistically significant number of measurements on which to base the mathematical fit.

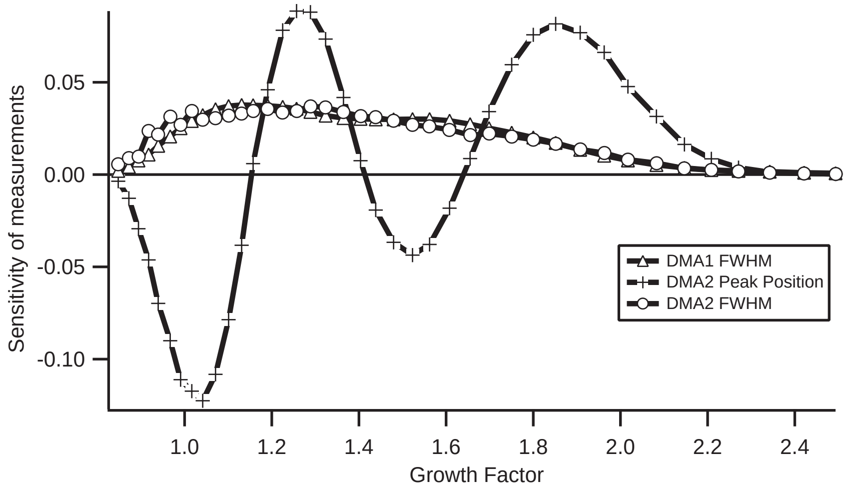

Fig. 5. The sensitivity of the growth factor measurements in Fig. 6 to parameters in the forward model. All three parameters have an effect, but the peak position gives the largest source of uncertainty.

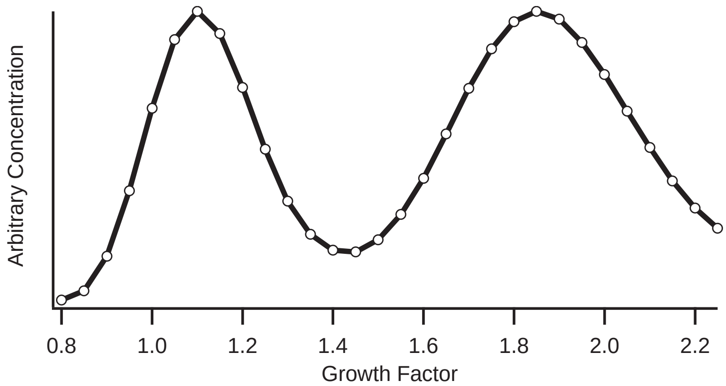

Fig. 6. The synthetic bi-modal measurement distribution.

9. Testing the retrieval process

In order to test the retrieval process, several synthetic measurement distributions were employed, and where appropriate results were compared with the gaussian fitting routine outlined in Section 4.

Firstly, a bi-modal measurement distribution was taken, creating two peaks using the same mathematical function as employed in the forward model. The measurement distribution is shown in Fig. 6. A successful retrieval should eliminate the transfer function from this distribution, retrieving a growth factor distribution consisting simply of two delta-functions. The retrieved distribution is shown in Fig. 7, showing the expected contributions into separated retrieval bins. The contributions are reported as a normalised probability, $\mathrm{d}C/\mathrm{d}g$, defined as the probability of a particle to fall into each of the retrieval bins (totaling 1), normalised to the area under the curve created by linear interpolation between the retrieved points.

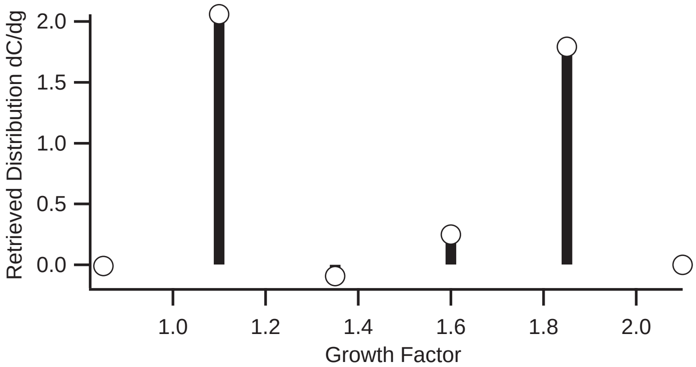

Fig. 7. The synthetic bi-modal retrieved distribution. The measurement distribution is retrieved into two separated retrieval bins.

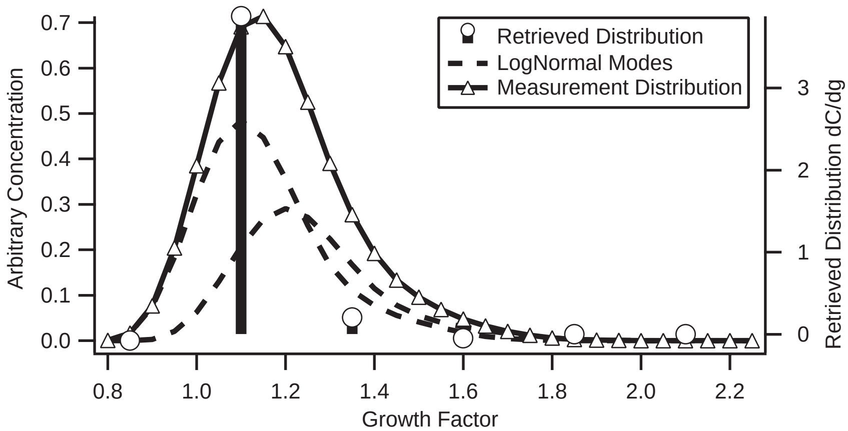

Fig. 8. The synthetic tailed distribution, comprised of two modes split by 0.125 in growth factor space. On retrieval the two modes both fall into the same bin.

To demonstrate the differences between the gaussian fitting approach and the OEM, a further synthetic measurement distribution was created, taking two lognormal modes split by 0.125 in growth factor space. In the retrieved distribution, shown together with the measurements in Fig. 8, the contributions from both modes fall into the same retrieval bin, as this degree of separation between the two modes cannot be resolved. This is emphasised by considering the averaging kernels in Fig. 9, whose half-widths indicate that the resolution of the retrieval is approximately 0.22 in growth factor space. This implies the retrieved distribution has been retrieved at the optimum resolution, and that within the errors in the system we are unable to resolve the two modes.

A further synthetic profile was then constructed, this time separating the two lognormal modes by 0.25 in growth factor space. The averaging kernels for the retrieval of this profile remain virtually identical to those shown in Fig. 9, indicating the optimum resolution of the retrieval is still approximately 0.22 in growth factor space. However, as shown in Fig. 10, as the two modes are now split by a factor at least as

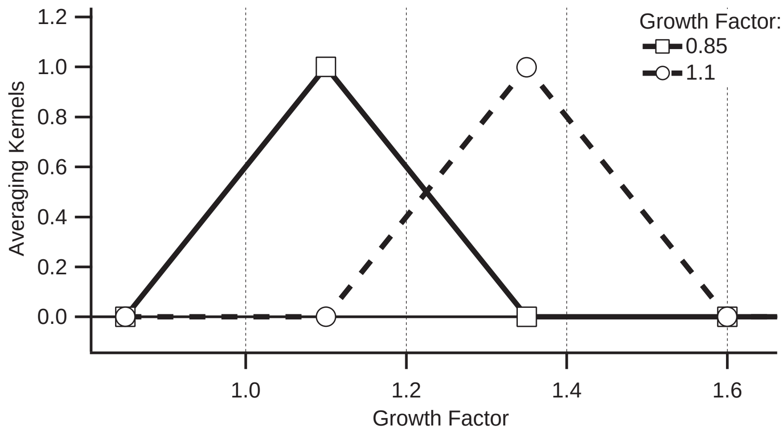

Fig. 9. Averaging kernels for the synthetic tailed distribution. The FWHM of the kernels indicate the resolution of the retrieval is approx 0.22 in growth factor space.

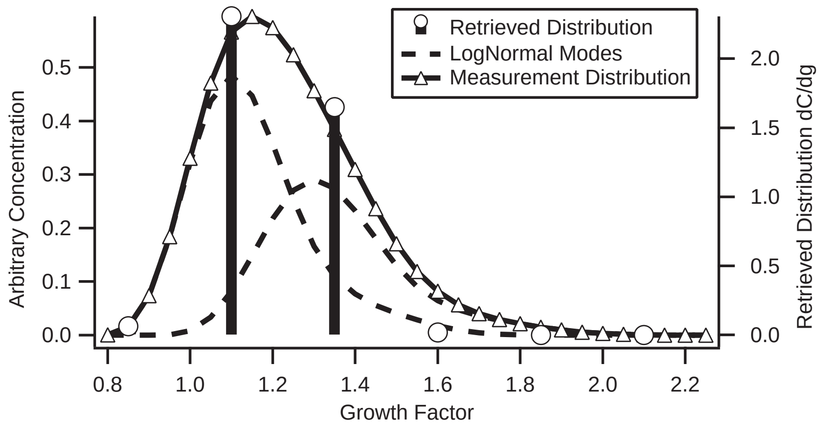

Fig. 10. The synthetic tailed distribution with larger separation (0.25 in growth factor space) between the two modes. The averaging kernels remain almost identical to Fig. 9, indicating the resolution is still 0.22. As can be seen the two modes can thus be resolved.

large as the resolution of the retrieval, their contributions now fall into two separate retrieval bins—i.e. we can resolve the two modes.

Finally, we consider the situation of the two modes split by 0.125 in growth factor space, but ignore the information presented by the averaging kernels and attempt to retrieve at a finer resolution so as to try and separate the modes. Fig. 11 shows the retrieved distribution and the averaging kernels in this case, and it is clear that the retrieval has failed. The distribution is now oscillatory in growth factor space, and the averaging kernels still indicate that the resolution of the retrieval is larger than the separation of the modes. It is thus not possible to resolve the modes simply by retrieving at a higher resolution.

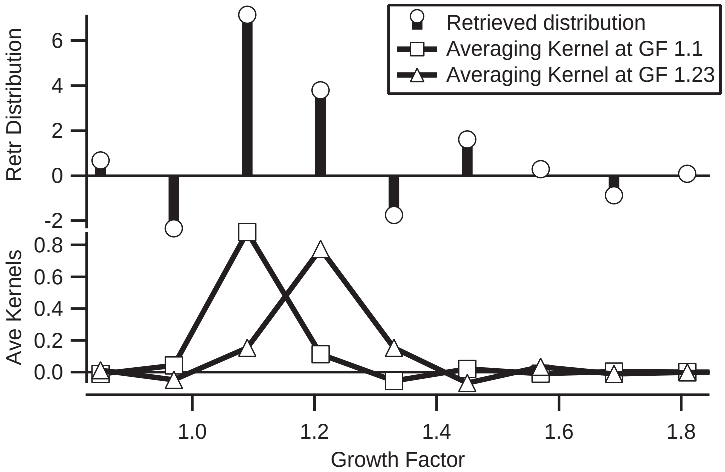

Fig. 11. Retrieved distribution and averaging kernels arising from trying to retrieve the original synthetic tailed distribution (separation 0.125) at a higher resolution. The retrieval and averaging kernels show oscillatory behaviour—i.e. the retrieval has failed. The kernels still indicate a resolution of approximately 0.2 in growth factor space.

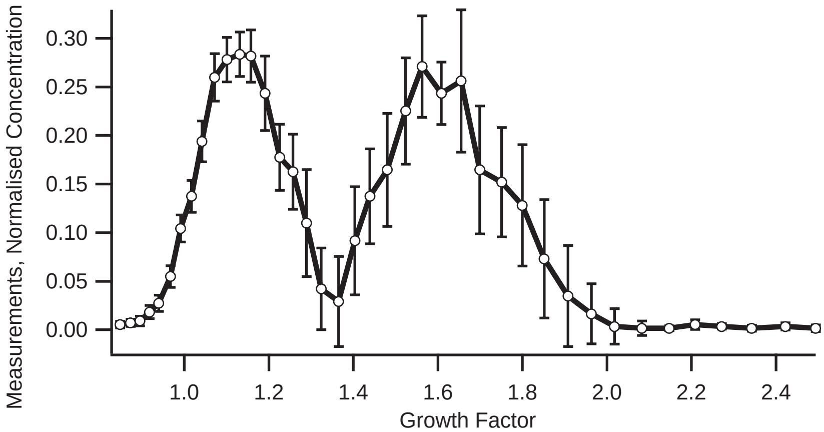

Fig. 12. The measurement distribution, **y**, and associated errors used in the example retrieval.

10. Example retrieval

An example measurement distribution, taken during the TORCH experimental campaign in Chelmsford, UK, during the summer of 2003, was chosen to fully test the retrieval process, including the sensitivity of the retrieval on the a priori information, $\mathbf{x_a}$, the errors in the a priori distribution, $\mathbf{S_a}$, and the errors in the measurements, $\mathbf{S}_\varepsilon$.

Fig. 12 shows the measurement distribution used, the error bars indicating the variances, i.e. the diagonal values of the measurement covariance matrix $\mathbf{S}_\varepsilon$. The errors used here are the calculated values for the instrument setup as it was during the time when the measurement distribution was taken. The longer error bars at the higher growth factors indicate that the RH errors discussed in the previous section are important in this particular measurement distribution.

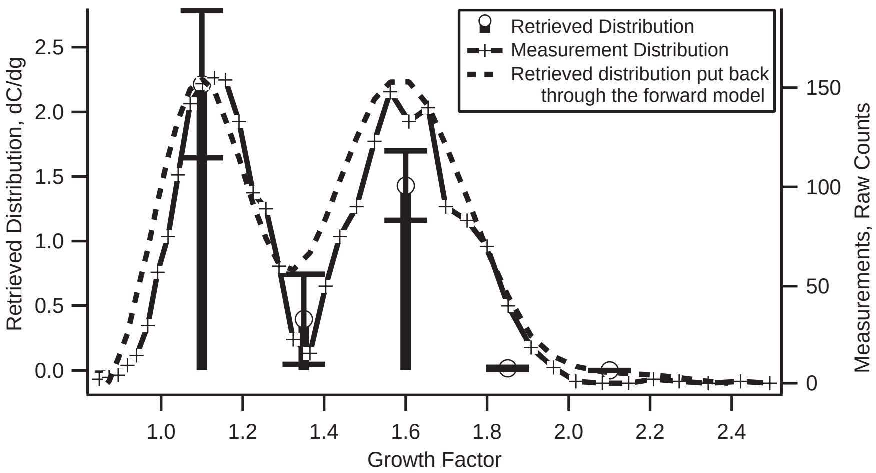

Fig. 13. The retrieved distribution, $\hat{\mathbf{x}}$, using resolution suggested by the averaging kernels. This distribution is run back through the forward model to verify the retrieval procedure.

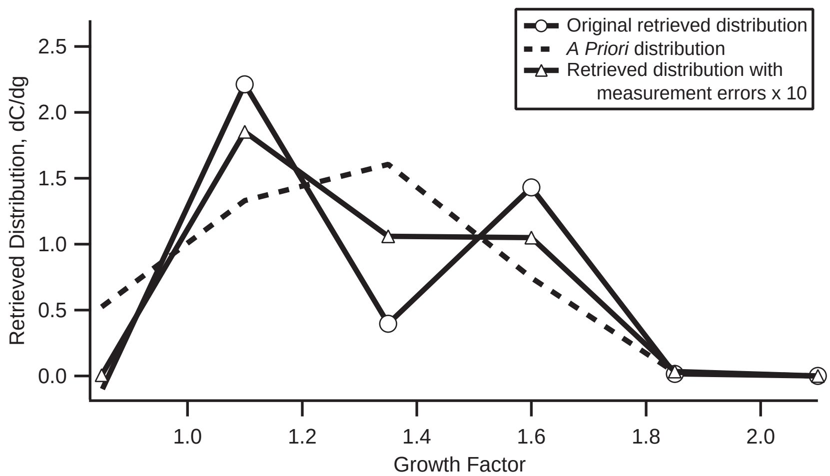

Fig. 14. The effect of increasing the measurements errors by a factor of 10. The retrieved distribution is biased towards the shape of the a priori distribution.

In Fig. 13 the retrieved distribution is shown, together with the result of running the retrieved distribution back through the forward model discussed in Section 6. It is clear that, within the errors, a successful retrieval should match the measurement distribution when operated on by the forward model. The error bars on the retrieved distribution are computed as described in Section 8, and it can be seen that, within errors, this line matches the measurement distribution, indicating a successful retrieval.

This retrieval has been made using the ideal resolution suggested by the averaging kernels as discussed in Section 8, suggesting that the resolution of the retrieval is approximately 0.24 in growth factor space.

The sensitivity of the retrieved distribution to the a priori distribution can be tested by adjusting the magnitude of the errors on the measurements and/or the a priori distribution. Fig. 14 shows the effect of

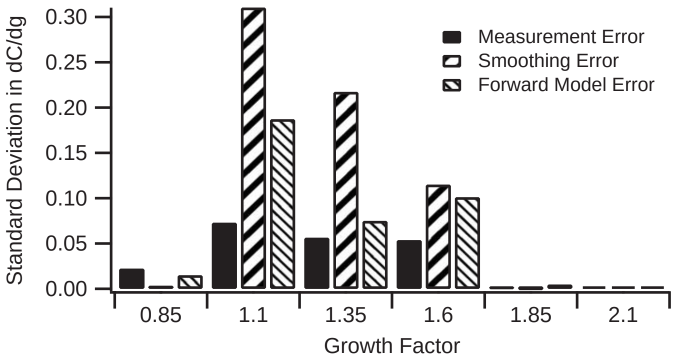

Fig. 15. The contributions of the smoothing error, the forward model error and the uncertainties in the measurements to the overall error in the example retrieval. This is a cumulative plot, the total height of the bars represents the overall error.

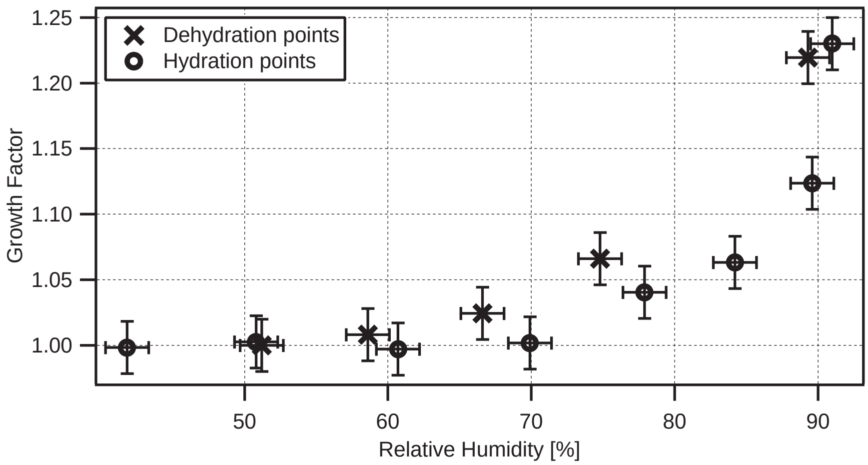

Fig. 16. A humidogram taken for 27 nm dry size particles in Manchester city centre on the 15th September 2003. The different curves for the hydration and dehydration of the aerosol can be ascertained.

increasing the measurement errors by a factor of 10. The shape of the retrieved distribution is skewed towards that of the a priori as the weighted average becomes biased towards the a priori distribution.

The contributions of each of the three components of the forward model to the overall error in the example retrieval are shown in Fig. 15. The smoothing effect is shown to make the most significant contribution to the overall total error, although the contributions from the measurement and forward model error are also significant.

11. Conclusions

A new design of hygroscopicity TDMA has been presented, with the capability of fast and accurate relative humidity control, allowing rapid, automatic determination of humidogram plots during ambient sampling. An example of this type of measurement is shown in Fig. 16. This method of data measurement

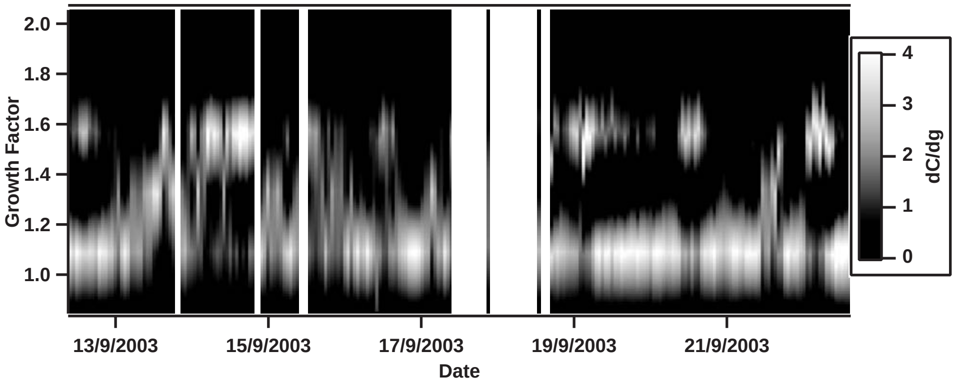

Fig. 17. Time series plot showing the retrieved growth factor distribution (at 90% RH) of 89 nm dry diameter particles sampled at an urban site in city centre Manchester, UK, during September 2003. The well-observed splitting of the aerosol population into less- and more-hygroscopic modes is observed. The small gaps in the time series represent periods when the instrument was taking humidogram measurements.

complements the in situ measurement of time series plots of the aerosol hygroscopic growth at 90% RH at a range of different dry diameters, an example of which is shown in Fig. 17. This figure shows the retrieved growth factor distributions as calculated by the OEM data analysis method.

The OEM method of data analysis has been shown to be a robust, thorough technique for the retrieval of TDMA data, giving an explicit error analysis and a measure of the resolution of the retrieval. It allows a comprehensive analysis of TDMA data, and the growth factor distribution retrieved is a significant improvement on determination of the modal peaks when considering closure studies with size-resolved compositional data. The OEM technique described here could also be applied to other instruments, most notably the volatility TDMA.

References

Cocker, D., Flagan, R., & Seinfeld, J. (2001). State of the art chamber facility for studying atmospheric aerosol chemistry. *Environmental Science and Technology*, *35*, 2594–2601.

Covert, D., & Heintzenberg, J. (1993). Size distributions and chemical properties of aerosol at Ny-Alesund, Svalbard. *Atmospheric Environment*, *A27*, 2989–2997.

Fuzzi, S., Laj, P., Ricci, L., Orsi, G., Heintzenberg, J., Wendisch, M., Yuskiewicz, B., Mertes, S., Orsini, S., Scwanz, M., Wiedensohler, A., Stratmann, F., Berg, O., Swietlicki, E., Frank, G., Martinsson, B., Gunther, A., Dierssen, J., Schell, D., Jaeschke, W., Berner, A., Dusek, U., Galambos, Z., Kruisz, C., Nesfin, S., Wobrock, W., Arends, B., & Ten Brink, H. (1998). Overview of the Po Valley fog experiment 1994 (CHEMDROP). *Contributions to Atmospheric Physics/Beitraege zur Physics Atmosphere*, *71*, 3–19.

Higgins, K. J., Jung, H., Kittelson, D. B., Roberts, J. T., & Zachariah, M. R. (2002). Size-selected nanoparticle chemistry: kinetics of soot oxidation. *Journal of Physical Chemistry A*, *106*, 96–103.

Maßling, A., Wiedensohler, A., & Busch, B. (1999). Concept of an advanced hygroscopic tandem differential mobility analyser with a great operational stability. *Journal of Aerosol Science*, *30*, S395–S396.

Maßling, A., Wiedensohler, A., & Voutilainen, A. (2003). Hygroscopic properties of submicrometer aerosol particles during ACE-Asia. In *Abstracts of the European aerosol conference*, 2003 (Vol. 1. pp. S9–S10).

McMurry, P. H., & Stolzenburg, M. R. (1989). On the sensitivity of particle size to relative humidity for Los Angeles aerosols. *Atmospheric Environment*, *23*, 497–507.

Pitchford, M., & McMurry, P. (1994). Relationship between measured water vapor growth and chemistry of atmospheric aerosol for Grand Canyon, Arizona, in winter. *Atmospheric Environment*, *28*, 827–839.

Rader, D., & McMurry, P. (1986). Application of the tandem differential mobility analyser for studies of droplet growth or evaporation. *Journal of Aerosol Science*, *17*, 771–787.

Rodgers, C. (1976). Retrieval of atmospheric temperature and composition from remote measurements of thermal radiation. *Reviews of Geophysics*, *14*, 609–624.

Rodgers, C. (1990). Characterization and error analysis of profiles retrieved from remote sounding measurements. *Journal of Geophysical Research*, *95*, 5587–5595.

Rodgers, C. (2000). Inverse methods for atmospheric sounding. *Theory and practice. Series on atmospheric*, *oceanic and planetary physics* (Vol. 2). New Jersey: World Scientific.

Svenningsson, B., Hansson, H.-C., Wiedensohler, A., Ogren, J. A., Noone, K. J., & Hallberg, A. (1992). Hygroscopic growth of aerosol particles in the Po Valley. *Tellus B*, *44*, 556–569.

Svenningsson, B., Hansson, H.-C., Wiedensohler, A., Noone, K. J., Ogren, J. A., Hallberg, A., & Colville, R. (1994). Hygroscopic growth of aerosol particles and its influence on nucleation scavenging in-cloud—experimental results from Kleiner–Feldberg. *Journal of Atmospheric Chemistry*, *19*(1–2), 129–152.

Svenningsson, B., Hansson, H., Martinsson, B., Wiedensohler, A., Swietlicki, E., Cederfelt, S., Wendisch, M., Bower, K., Choularton, T., & Colvile, R. (1997). A closure study of sub-micrometer aerosol particle hygroscopic behaviour. *Atmospheric Environment*, *31*, 2463–2475.

Swietlicki, E., Zhou, J., Berg, O., Martinsson, B., Frank, G., Cederfelt, S., Dusek, U., Berner, A., Birmili, W., Wiedensohler, A., Yuskiewicz, B., & Bower, K. (1999). A closure study of sub-micrometer aerosol particle hygroscopic behaviour. *Atmospheric Research*, *50*, 205–240.

Vakeva, M., Kulmala, M., Stratmann, F., & Hameri, K. (2001). Field measurements of hygroscopic properties and state of mixing of nucleation mode particles. *Atmospheric Chemistry and Physics Discussion*, *1*, 409–410.

Weingartner, E., Gysel, M., & Baltensperger, U. (2002). Hygroscopicity of aerosol particles at low temperatures. 1. New low-temperature H-TDMA instrument: setup and first applications. *Environmental Science and Technology*, *36*, 55–62.

Winkelmayr, W., Reischl, G., Lindner, A., & Berner, A. (1991). A new electromobility spectrometer for the measurement of aerosol size distributions in the range from 1 to 1000 nm. *Journal of Aerosol Science*, *22*, 289–296.

Zhang, X., McMurry, P., Hering, S., & Casuccio, G. (1993). Mixing characteristics and water content of submicron aerosols measured in Los Angeles and at the Grand Canyon. *Atmospheric Environment A*, *27*, 1593–1607.

Zhou, J., Swietlicki, E., Hansson, H., & Artaxo, P. (2002). Sub-micrometer aerosol particle size distribution and hygroscopic growth measured in the Amazon rain forest during the wet season. *Journal of Geophysical Research*, *107*.