Pergamon

Atmospheric Environment Vol. 30, No. 6, pp. 929–940, 1996

1352-2310/96 $15.00 + 0.00

1352-2310(95)00295-2

NOVEL APPROACH FOR TOMOGRAPHIC RECONSTRUCTION OF GAS CONCENTRATION DISTRIBUTIONS IN AIR: USE OF SMOOTH BASIS FUNCTIONS AND SIMULATED ANNEALING

A. C. DRESCHER,*† A. J. GADGIL,† P. N. PRICE† and W. W. NAZAROFF*‡

*Civil and Environmental Engineering Department, University of California, Berkeley, CA 94720-1710, U.S.A.; †Indoor Environment Program, Energy and Environment Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, U.S.A.

(*First received* 3 *January* 1995 *and in final form* 25 *July* 1995)

Abstract—Optical remote sensing and iterative computed tomography (CT) can be applied to measure the spatial distribution of gaseous pollutant concentrations. We conducted chamber experiments to test this combination of techniques using an open path Fourier transform infrared spectrometer (OP-FTIR) and a standard algebraic reconstruction technique (ART). Although ART converged to solutions that showed excellent agreement with the measured ray-integral concentrations, the solutions were inconsistent with simultaneously gathered point-sample concentration measurements. A new CT method was developed that combines (1) the superposition of bivariate Gaussians to represent the concentration distribution and (2) a simulated annealing minimization routine to find the parameters of the Gaussian basis functions that result in the best fit to the ray-integral concentration data. This method, named smooth basis function minimization (SBFM), generated reconstructions that agreed well, both qualitatively and quantitatively, with the concentration profiles generated from point sampling. We present an analysis of two sets of experimental data that compares the performance of ART and SBFM. We conclude that SBFM is a superior CT reconstruction method for practical indoor and outdoor air monitoring applications.

Key word index: Remote sensing, computed tomography, simulated annealing, gas monitoring.

INTRODUCTION

Measuring the spatial distribution of airborne gas and particle concentrations is an important objective in many contexts, including fundamental studies of flame dynamics, investigations of pollutant dispersion from localized sources in industrial hygiene settings, and characterizing the concentrations of air toxics near an industrial facility. Measurement programs based on sampling at discrete points encounter significant limitations, such as poor temporal and spatial resolution, intrusiveness, and equipment and labor intensiveness. To overcome some of these problems, remote sensing coupled with computed tomography (CT) has been suggested as an alternative means to measure the distribution of contaminant concentrations in air (Wolfe, 1980). In studies of flame dynamics, this approach is already routinely applied (Faris and Byer, 1987). The use of iterative CT methods in conjunction with one or more open path Fourier transform infrared (OP-FTIR) spectrometers has been proposed for monitoring harmful gases in indoor workplaces (Todd and Leith, 1990; Todd, 1992a, b). Potential applications to monitoring emissions from outdoor air sources have also been explored (Fischer *et al.*, 1994).

Yost *et al.* (1994) reported the first test of indoor air concentration measurements in a horizontal plane using the CT-FTIR method in a chamber. The test used a scanning geometry comprising four sets of parallel rays, with each set rotated 45° with respect to the previous one. Although ideal for iterative CT reconstruction, this approach suffered from limited temporal resolution. Only one source and detector were available for those experiments. Both had to be moved substantially to produce the desired ray pattern. The resulting scan time was 3 h. During this time, fluctuations in the concentration pattern, and possibly changes in the air flow and dispersion due to the variable positioning of the equipment, produced inconsistent ray-integral concentration measurements. Spatial resolution was also limited due to the relatively large distance between parallel rays. Thus, the CT reconstruction in this study was not entirely satisfactory. While the peak concentration was correctly located, most of the ray-integral concentrations

‡Author to whom correspondence should be addressed. Fax: (510) 642–7483, E-mail: nazaroff@ce.berkeley. edn.

of the reconstructed profile did not agree well with the experimental ray-integral data.

During October and November 1993, we undertook a series of experiments designed to improve the CT-FTIR method by using a steerable FTIR instrument in conjunction with fixed mirrors and retroreflectors at the perimeter of a test chamber. With this system, the total time for a single scan was reduced from 3 h to just over 6 min.

By necessity, the ray geometry differed substantially from that employed in the experiments of Yost *et al.* (1994). This new ray geometry, though permitting us to overcome largely the problem of temporal variability, created false spatial variability in the reconstructed concentration profiles when used in conjunction with the standard algebraic reconstruction technique (ART) suggested by Todd and Leith (1990) and Todd and Ramachandran (1994). Because the rays in the new geometry were less "independent" than would be an equal number of rays in parallel projections, they produced a more underdetermined system of equations when using the number of pixels necessary to accurately capture spatial variability in the concentration profiles.

Using two of the experimental runs as illustrative examples, this paper describes how CT reconstructions were first attempted with ART, how and why they failed, and then introduces a new reconstruction method that overcomes the difficulties of ART. The new method is based on modeling the concentration distribution with smooth basis functions (using bivariate Gaussians in this case). For the two experimental runs considered, the paper compares measured point concentration data, concentration profiles generated by ART and concentration profiles generated by the new smooth basis function minimization technique (SBFM). The full set of experiments is described in detail in a related paper (Drescher *et al.*, 1995).

METHOD

Experiments

Experiments were conducted in a $6.4 \times 7.3 \times 2.4$ m high chamber, comprising a wooden frame covered by a thick polyethylene liner, located within a larger building. One short wall had two 46 cm high $\times$ 6.4 m openings centered at heights of 67 cm and 1.75 m. Along the opposite side of the chamber, in the ceiling, was a 92 cm wide slot. All openings were covered with coarse-fibered filter material. The ceiling slot was connected to a hood through which air was drawn from the chamber at a rate of $\sim 5600\ m^3\ h^{-1}$, i.e. approximately 50 room volumes per hour.

In each experiment, a buoyant mixture of 10% sulfur hexafluoride (SF_6) in helium was released at a constant rate from a 2.5-cm diameter, porous ceramic sphere, placed 60 cm above the floor. Once steady-state conditions were established, SF_6 concentrations were measured over a period of 1 h in a horizontal plane 1.3 m above the floor. Two measurement methods were employed and these will be referred to as "point-sample" measurements and "ray-integral" or "path-integral" measurements. (1) Air samples were collected continuously into gas sampling bags at 26 points and were subsequently analyzed for SF_6 by a photoacoustic detector (Bruel & Kjaer [B&K], Model 7620). During each experiment, the photoacoustic detector was used to sample air from five additional points in the plane (and one point in the fume hood, sampling the chamber exhaust air). Air from each of these points was drawn continuously, at a flow rate of $1\ \ell\ min^{-1}$, through a 2.3 ℓ glass jar and sampled from the jar by the detector once every 6 min. For the present analysis, SF_6 concentrations at these five points were time-averaged over the duration of the experiment. (2) Path-integral SF_6 concentrations were measured over 56 intersecting ray paths using a centrally placed, steerable OP-FTIR system, four large retroreflectors and 52 flat mirrors placed along the perimeter of the chamber. (See Fig. 1.) Nine complete sweeps (of all 56 rays) were gathered over the course of 1 h for each experiment.

Data processing

For each orientation of the FTIR, the mean, μ_r, and variance, σ_r^2, of the ray-integral concentrations for the nine measurements were computed. Half of the CT reconstructions were based on the measured time-average ray-integral concentrations, μ_r. The other half were based on synthetic, noise-free ray-integral concentration data.

Figure 2 summarizes how the point-sample data and the ray-integral data were processed. Six maps of SF_6 concentration as a function of position in the plane, designated by letters (a)–(f), were generated for each experiment. For the first five maps, the measurement plane was divided into 195 (13×15) square pixels (0.22 m^2 each). The point sample data (a) were kriged (Jones *et al.*, 1986) using the graphing program SURFER (Golden Software, Golden, CO) and the value at the center of each pixel was plotted in graphs marked (b). The kriged point sample data were also used to generate "synthetic ray integrals" (strip sums, based on the assumption of a uniform concentration in each pixel) for the 56 ray paths used in the experiment. These synthetic ray-integral concentrations were entered into the ART-CT program to test the algorithm under noise-free conditions and compare the resulting concentration maps, marked (c), with those based on the measured ray-integral concentrations, marked (d). SBFM-CT reconstructions were also performed on the synthetic (e) and the measured (f) ray-integral data.

A figure of merit was defined to quantify the overall difference between the various CT reconstructions and the measured point-sample data:

$$\sigma_{pnt} = \left(\frac{\sum_{pt=1}^{31} (C_{pt,m} - C_{pt,CT})^2}{\sum_{pt=1}^{31} (C_{pt,m})^2} \right)^{1/2} \quad (1)$$

where $C_{pt,m}$ is the measured point-sample concentration at a given location and $C_{pt,CT}$ is the reconstructed concentration at that location. A value of zero for the parameter σ_{pnt} would indicate perfect agreement between the measured point-sample data and the concentrations at those same points calculated by the CT programs.

An equivalent parameter was defined for the CT reconstructions based on synthetic ray integrals calculated from the kriged point concentration data:

$$\sigma_{pix} = \left(\frac{\sum_{p=1}^{195} (C_{p,k} - C_{p,CT})^2}{\sum_{p=1}^{195} (C_{p,k})^2} \right)^{1/2} \quad (2)$$

where $C_{p,k}$ is the concentration in a pixel obtained from kriging and $C_{p,CT}$ is the pixel concentration calculated by the CT reconstruction program.

POINT SAMPLE DATA

Figures 3a and 4a show the time-averaged SF_6 concentrations at the sampling points in the plane of measurement for Experiments 1 and 2, respectively.

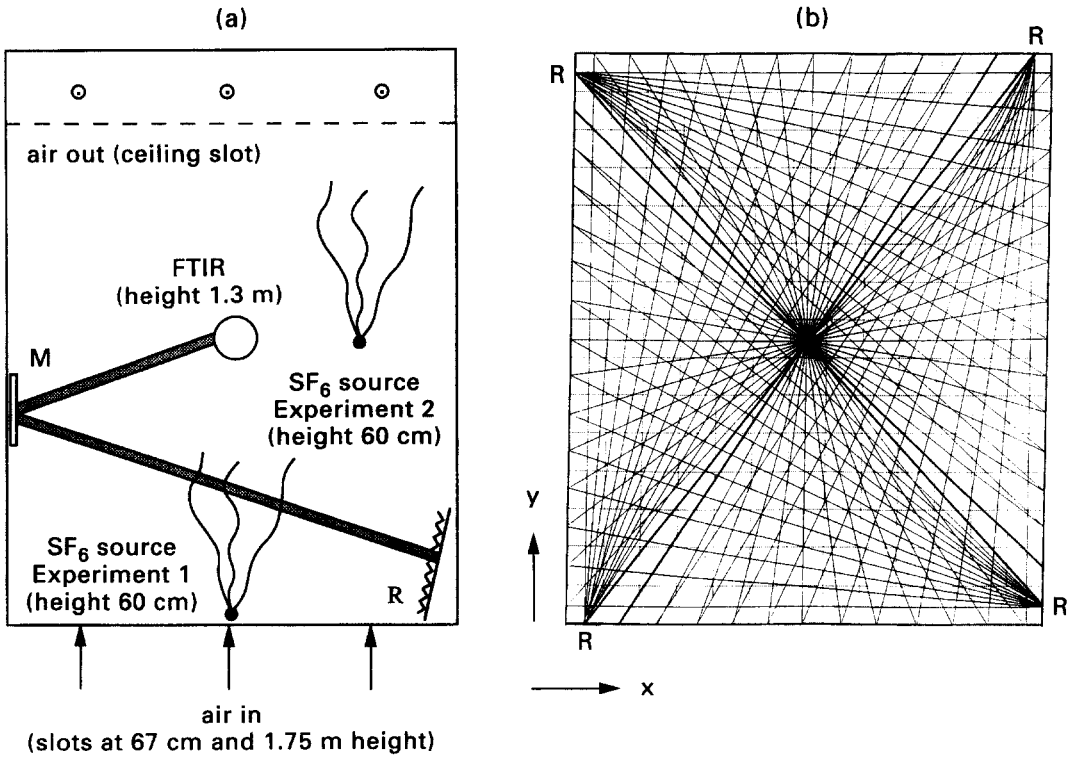

Fig. 1. (a) Top view schematic of the experimental chamber: M = mirror and R = retroreflector. For clarity, only one mirror and one retroreflector are shown. (b) Top view of the measurement plane, showing pixels (dotted lines), and rays (solid lines) emanating from the centrally located FTIR (with coincident source and detector). Rays in bold are reflected directly back to the center; other rays are first reflected off of a mirror along the chamber perimeter and then off of a retroreflector (R). Pixels are 47 cm on a side. Ray diameter is approximately 20 cm.

The location of the SF_6 source and the general direction of air flow in each case is also shown. It appears that in Experiment 1, one major and one minor concentration peak were present, while for Experiment 2 only one major peak can clearly be distinguished. This interpretation is consistent with the results of smoke-tube tracer tests we performed. Smoke released near the SF_6 source point in Experiment 1 was swept immediately up through the measurement plane by the incoming air stream and was dispersed further downstream. This transport pattern would be expected to produce a high concentration of SF_6 just downstream of the source, and a secondary, broader peak in concentration further downstream, as the SF_6 is dispersed back down into the measurement plane. By contrast, in Experiment 2, smoke released near the SF_6 source was immediately dispersed significantly in all directions while slowly being advected in the general direction of the exhaust hood.

Kriged point-sample data are presented in Figs 3b and 4b. The broadness of the dominant peak in Fig. 3b and the local minimum near position ($x = 5.5$, $y = 13.5$) in Fig. 4b are artifacts of the kriging interpolation. There is no evidence for these features in the point-sample data. This point should be kept in mind when comparing these figures with the CT reconstructions of the experimental ray-integral concentration data.

ART RECONSTRUCTION

Method

The ART algorithm iteratively adjusts concentrations in pixels along each ray in sequence, so that the difference between the measured ray-integral concentration and that calculated using the inferred pixel concentrations is minimized (Brooks and DiChiro, 1975). ART effectively weights all rays equally in the reconstruction process, as opposed to, for instance, weighting rays by their measurement uncertainty. In our implementation, the initial guess for the concentration distribution was a uniform distribution. We first obtained a ray-average concentration for each ray by dividing the ray-integral concentration by the area of the ray in the measurement plane. The mean of all 56 ray-average concentrations was then taken to be the first guess for the concentration in each pixel.

After each iteration, parameter σ_{ray} (defined in analogy with σ_{pnt} and σ_{pix}) was evaluated:

$$\sigma_{ray} = \left(\frac{\sum_{r=1}^{56} (PIC_r - PIC_{r,CT})^2}{\sum_{r=1}^{56} (PIC_r)^2} \right)^{1/2} \quad (3)$$

where PIC_r is the "true" ray-integral concentration for ray r ($PIC_r = \mu_r$ for FTIR data, or PIC_r = calculated value for kriged point concentration data), and

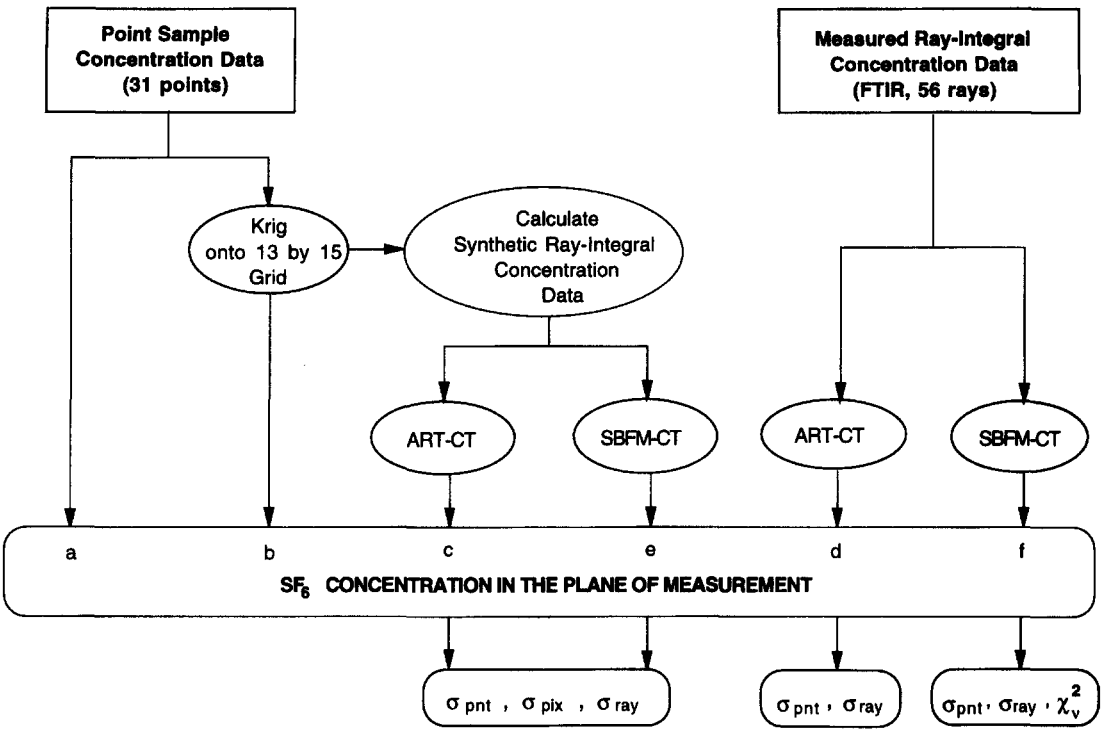

Fig. 2. Summary of data processing approach. Rectangles represent experimental input data; ellipses represent calculations; rectangles with rounded corners represent results. Lower case alphabets in the long horizontal rectangle are keyed to results shown in Figs 3 and 4.

$\mathrm{PIC}_{r,\mathrm{CT}}$ is the path-integral concentration calculated from the reconstructed profile. For ART,

$$\mathrm{PIC}_{r,\mathrm{CT}} = \sum_{p=1}^{195} a_{p,r} C_{p,\mathrm{CT}} \qquad (4)$$

where $a_{p,r}$ is the area of pixel p intercepted by ray r.

The parameter σ_{ray} quantifies the overall agreement between the "true" ray-integral concentrations, PIC_r, and those calculated from the reconstructed concentrations. For the kriged data, the ray integral, PIC_r, was calculated as a strip sum using the experimental ray width as the width of the strip. The iteration process was terminated either when the fractional change in σ_{ray} from one iteration to the next was less than 0.5% or when the value of σ_{ray} fell below 0.001.

The order in which ART was applied to rays to reevaluate the pixel concentrations along them was found not to influence the final solution. The only constraint applied to the reconstruction process was that pixel concentrations could not be less than zero.

Results

Figures 3c and 4c show the concentration profiles produced by applying ART to the synthetic ray-integral data generated from the kriged point-sample measurements. The very low values of σ_{ray} (see Table 1), show that the reconstructions generated by ART are valid solutions given the synthetic ray-integral concentrations. In fact, the value of σ_{ray} can be made as low as desired (without changing the qualitative features of the concentration map) by increasing the number of iterations. However, since the input data for these reconstructions derive from a known concentration profile and were noise free, one would expect a perfect match between frames (b) and (c). In other words, Figs 3b and 4b each represent a solution to the respective synthetic data inversion problem solved by ART. It is apparent from Figs 3c and 4c that ART did not converge to these "true" solutions. The peak of the concentration is correctly located in each case, but the peak magnitude is incorrect. In addition, many secondary artifacts not present in the original kriged maps, such as minor peaks and ridges, appear in the ART reconstructions. The overall discrepancy is manifested in the relatively high values of σ_{pix} shown in the first column of Table 1: 0.62 for Experiment 1 and 0.45 for Experiment 2.

By starting the ART reconstruction with different first guesses or by manipulating the algorithm in other ways the program could be made to converge to other solutions which are equally valid but distinct in detail. Only if the first guess was very close to the true solution did ART converge to the true solution. These unsatisfactory features are a consequence of a basic underlying problem: ART embodies a seriously underdetermined system. With 195 unknowns (pixel concentrations) and only 56 equations (ray integrals), problems with indeterminacy were to be expected. Previous simulation work (Todd and Leith, 1990; Todd and Ramachandran, 1994) has shown that,

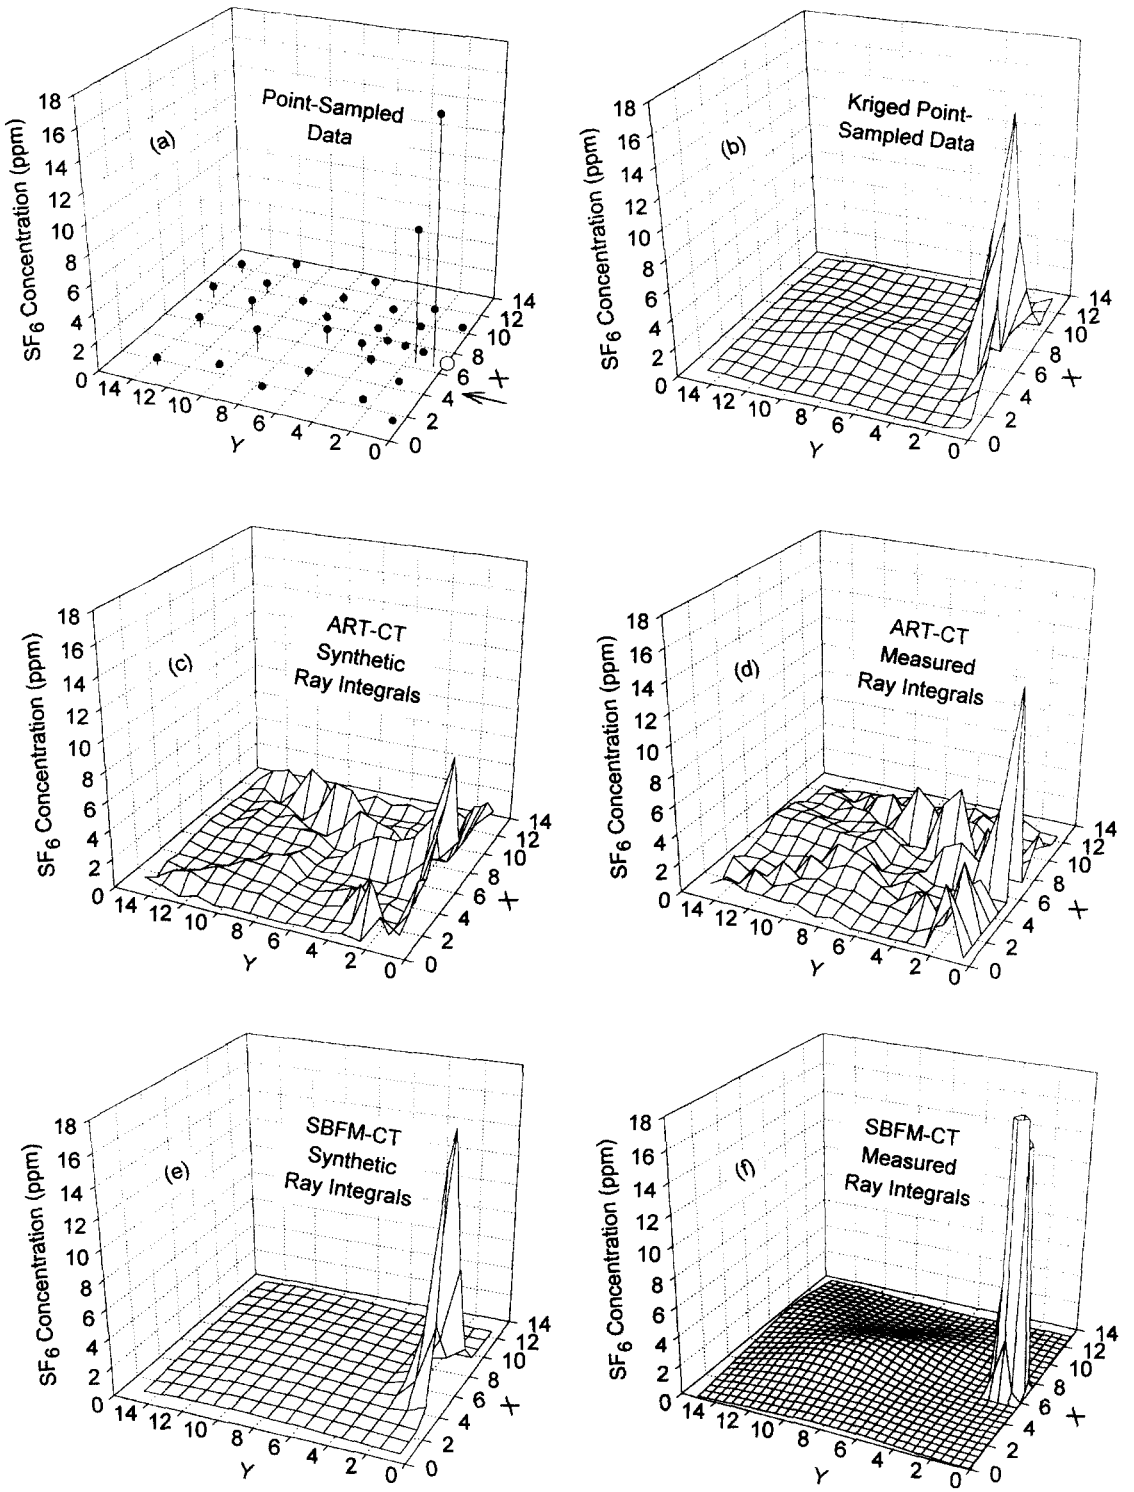

Fig. 3. Results for Experiment 1. The x and y axes are in pixel units. (a) Point sample data (source shown as a white dot). The arrow indicates the direction of the ventilation airflow. (b) Kriged point sample data. (c) ART reconstruction using synthetic ray-integral data. (d) ART reconstruction using experimental ray-integral data. (e) SBFM reconstruction using synthetic ray-integral data (2 Gaussians). (f) SBFM reconstruction using experimental ray-integral data (2 Gaussians). Only concentrations up to 18 ppm are shown, so that the low concentrations can be distinguished. The peak concentration of the narrow Gaussian in (f) is 43 ppm.

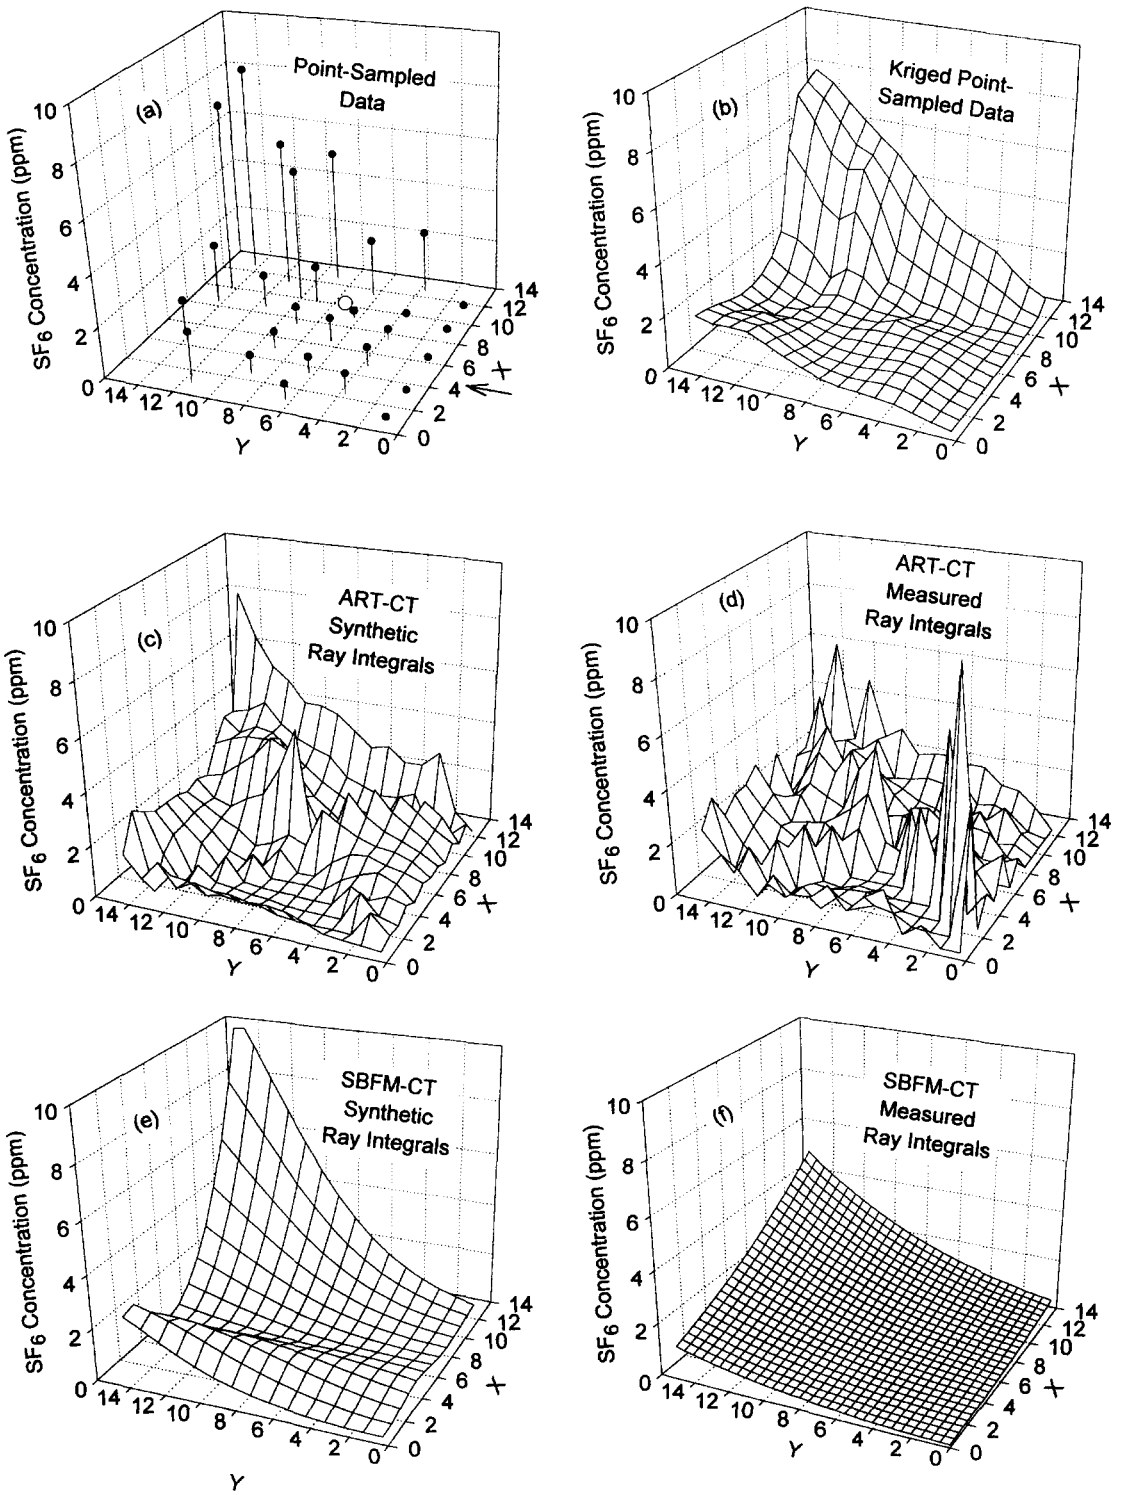

Fig. 4. Results for Experiment 2. The x and y axes are in pixel units. (a) Point sample data (source shown as a white dot). The arrow indicates the direction of the ventilation airflow. (b) Kriged point sample data. (c) ART reconstruction using synthetic ray-integral data. (d) ART reconstruction using experimental ray-integral data. (e) SBFM reconstruction using synthetic ray-integral data (2 Gaussians). (f) SBFM reconstruction using experimental ray-integral data (1 Gaussian). Peak concentration of 10.6 ppm clipped in e to show results on the same scale as a–d, f. During this experiment, three of the pumps used to draw air into the gas bags failed. Therefore, only 28 time-averaged point concentrations are shown in a.

under favorable conditions, good reconstructions can be obtained for an underdetermined system with the ratio of the number of pixels to the number of rays being as large as 6.7:1. In those reported cases, the reconstructed pixel concentrations matched well the true distribution. However, the concentration profiles used in the work of Todd and coworkers were less challenging (in the sense that they mitigated the problem of indeterminacy) because they were composed of narrow Gaussian peaks centered within the reconstruction plane. Complete, narrow, and distinct peaks pose a relatively easy challenge for ART, because only a small fraction of the rays pass through the peak, thus limiting the extent to which the detected species can be spatially dispersed in the reconstruction into so-called ghost images. In contrast, the broad peaks evident in our kriged experimental data are intercepted by almost all rays. This feature magnifies the indeterminacy problem by increasing the number of independent ways in which the species mass can be allocated amongst pixels without affecting the ray-integral concentrations.

The problem of indeterminacy can be mitigated by using more ideal ray geometries, such as parallel projections. However, if the scanning time is to remain short, this approach requires use of many sources and detectors or ample steerable optics. The cost of these options is prohibitive.

Another approach to lessen the impact of indeterminacy is to make individual pixels larger, thereby reducing the number of pixels and correspondingly decreasing the number of unknowns. However, the number of pixels is directly related to the spatial resolution of the reconstruction, so that a decrease in the number of pixels may reduce the resolution to the point where the true spatial variation in concentration can no longer be accurately represented. We attempted ART on the measured ray integral concentrations of Experiment 1 with a pixel grid of 7×8 (i.e. 56 pixels, equal to the number of rays). ART could not generate a reconstruction with a σ_{ray} below 0.3, indicating that *no* adequate solution to the ray integrals exists at this coarse discretization.

Figures 3d and 4d show the ART reconstructions of the measured (FTIR) ray-integral concentration data. These concentration maps are significantly less satisfactory than the ART reconstructions of the synthetic ray-integral data. Figure 4d shows the peak concentration for Experiment 2 to be near the origin. There is no evidence for this in the measured point data, and it is highly unlikely that this was the true location of the peak concentration given the position of the SF6 source and direction of air motion (as shown in Fig. 4a). Nevertheless, the goodness-of-fit parameter σ_{ray} is low (0.026) in this case, tempting one to believe that a correct reconstruction has been achieved. It is clear that applying ART to the reconstruction of gas concentration profiles in air using practical, albeit nonideal ray geometries can lead to serious misrepresentations.

SBFM RECONSTRUCTION

In view of these problems with ART, we developed an alternative approach that eliminates indeterminacy. This approach recognizes that the solutions can be constrained to a smaller domain of realistic possibilities for a given experimental context by imposing a functional form on the concentration profile. The artificial discretization of space into pixels is not necessary for reconstructions based on measured ray-integral data. The many sharp and mostly unrealistic gradients at pixel boundaries evident in the ART reconstructions are avoided. For gaseous species concentrations, one expects that turbulent air motion will tend to disperse pollutants and thus smooth out concentration gradients. With only a few pollutant sources, no more than three or four local concentration maxima would be anticipated under normal ventilation conditions.

Table 1. Summary of goodness-of-fit parameters for CT reconstructions

	Synthetic ray-integral data		Measured ray-integral data	
	ART	SBFM	ART	SBFM
Experiment 1				
σ_{pnt}	0.53	0.26	0.41	0.11
σ_{pix}	0.62	0.33	—	—
σ_{ray}	0.005	0.067	0.047	0.084
χ^2_ν	—	—	—	1.19
Figure	3c	3e	3d	3f
Experiment 2				
σ_{pnt}	0.49	0.25	0.93	0.51
σ_{pix}	0.45	0.21	—	—
σ_{ray}	0.002	0.04	0.026	0.16
χ^2_ν	—	—	—	0.11
Figure	4c	4e	4d	4f

Method

In the newly developed approach, the concentration distribution to be determined (from synthetic or measured ray-integral data) was modeled as a superposition of bivariate Gaussians defined in the x–y plane. The Gaussian was chosen as a basis function because it is easy to calculate and yet versatile enough to satisfactorily represent a broad range of possible concentration profiles. Each Gaussian has six adjustable parameters: x and y coordinates of the peak, standard deviations along its two principal axes, peak height (maximum concentration) and angle of rotation of the principal axes relative to the x and y axes.

The ideal Gaussian parameters were assumed to be those for which σ_{ray}, as defined in equation (3), was minimized. Following standard experimental data fitting theory (Bevington, 1969), an alternative objective of the reconstruction procedure might have been to minimize the *weighted* sum of squares of deviations of the data, PIC_r, from the ray sums calculated using the Gaussian concentration profiles, $PIC_{r,CT}$. Thus, for the measured ray-integral data, the reduced chi-squared, defined below, might have been minimized.

$$\chi_\nu^2 = \frac{1}{(56-n)} \sum_{r=1}^{56} \frac{1}{\sigma_r^2} (PIC_r - PIC_{r,CT})^2 \qquad (5)$$

where ν is the number of degrees of freedom after adjusting n fitting parameters to match the 56 ray integral concentration values ($\nu = 56 - n$ where n is six times the number of Gaussians used for the fit). However, after analyzing the two data sets presented here and the steady-state experiments described in Drescher *et al.* (1995), it was found that minimizing the reduced chi-squared often did not lead to the most accurate reconstruction, for the reasons outlined in the two paragraphs below.

For cases with a large number of degrees of freedom, the reduced chi-squared will be close to unity if (1) the correct fitting function is chosen and (2) the measurement errors, σ_r, are Gaussian distributed and accurately estimated. For a set of measurements of an unchanging quantity, the measurement uncertainty would be calculated as the error in the mean (σ_r divided by the square root of the number of observations, N). However, in our case, the measured quantity (the ray-integral concentration for each ray) was itself varying in time. The uncertainty was therefore estimated as σ_r. Consequently, the weighting factor for the deviation term for each ray-integral concentration is simply the inverse of its variance rather than N/σ_r^2.

Careful analysis of the data for the 11 steady-state experiments (nine of which are reported in Drescher *et al.* (1995)) revealed two important factors. (1) For many rays, especially those passing through areas of the plane near the SF_6 source, the deviations in the measured path-integral concentrations from the mean were not Gaussian distributed. Thus, one of the criteria for the proper calculation of χ_ν^2 was not satisfied. (2) Moreover, variability in path-integral concentrations of rays having a relatively low μ_r was often disproportionately low. Because the weighting factor for each ray in the calculation of χ_ν^2 was the standard deviation *squared*, σ_r^2, the minimization of χ_ν^2 tended to put undue emphasis on the rays with low μ_r. Thus, the peak concentrations in the reconstruction would tend to be much less accurate than when minimizing σ_{ray}. These two factors led us to conclude that, while calculating χ_ν^2 for a reconstructed profile provided useful information about the quality of the reconstruction, it was not the ideal function to be minimized in the SBFM approach.

For SBFM-CT reconstructions using measured ray integrals (profiles marked (f) in our case) the concept of pixels is superfluous and was discarded. Instead, each ray-integral concentration, $PIC_{r,CT}$ was first evaluated analytically as a line integral along the centerline of the ray. This result was then multiplied by the width of the ray to find the value of the strip integral. Given that the ray width was 43% of the width of the pixels employed in all other reconstructions, the spatial resolution was approximately doubled by this procedure. For this reason, the graphs shown in Figs 3f and 4f were plotted at twice the grid density compared to the other surface plots. Moreover, the resulting concentration distribution was no longer a discontinuous function, thus making it a less artificial representation of reality. In this lies the crux of the SBFM method: we impose a model on the underlying concentration distribution that is closer to reality than ART's model (completely independent concentration values distributed into pixels). The fact that our model is much more constraining results in a decrease in the number of degrees of freedom available to the reconstruction and thus a decrease in the number of possible solutions.

The SBFM program was written in Turbo Pascal and based on the Amebsa routine in Numerical Recipes (Press *et. al.*, 1992). It uses a combination of the simplex method and a simulated annealing technique to find the set of Gaussian parameter values that minimizes σ_{ray}. Simulated annealing can find the global minimum of a function (rather than a local minimum close to the first guess) by mimicking the gradual cooling of a crystallizing liquid. For such a liquid, the lowest energy state is that of the perfect crystal (global minimum) rather than that of the amorphous solid (local minimum) produced by rapid cooling.

One issue of practical importance is selecting an appropriate number of Gaussians for the reconstruction. For each experiment in the present study, the SBFM program was first run using only one fitting parameter: a constant concentration in the plane of reconstruction. Then the program was run using one Gaussian, two Gaussians and, for Experiment 1, three Gaussians. Each time a Gaussian was added, the parameters of the previously obtained peaks were

used as a starting condition to reduce the computation time. The values of χ_ν^2 or σ_{ray} were recorded as a function of the number of fitting parameters. When the relative improvement (decrease) in χ_ν^2 or σ_{ray} became small compared to the increase in the number of fitting parameters, it was assumed that a reasonable fit had been achieved. The synthetic data were modeled with two Gaussians in the SBFM approach, since this number of Gaussians seemed the minimum number necessary to be able to capture the major features apparent in the kriged maps of concentration.

Results

Figures 3e and 4e show the results of the reconstruction from synthetic ray-integral data by SBFM using two Gaussians. Clearly, the agreement with the "true" concentration profiles (Figs 3b and 4b, respectively) is superior to the ART reconstructions. This superiority is also evident in that the values of σ_{pix} (Table 1, second column) are approximately half of what they are for ART. On the other hand, the σ_{ray} values are more than ten times higher with SBFM than with ART, although still quite good at 0.067 and 0.040 for the two experiments. The poorer value of σ_{ray} under SBFM reflects the fact that there are fewer degrees of freedom for matching ray integrals relative to ART.

Compared with ART, the maps generated by applying SBFM to the experimental ray-integral concentration data (Figs 3f and 4f) are much more consistent with the measured point-sample data. This is also reflected in the significantly lower σ_{pnt} values of the SBFM reconstructions (see Table 1). The near-source differences between point-sample data and SBFM reconstructions may be real, since high spatial and temporal variability close to the SF_6 source, combined with the differences in sampling methods between the point-sample and ray-integral measurements would lead to somewhat different information contained in each data set. For instance, time variations are completely smoothed in the point-sample data, but they are a source of noise in the ray-integral reconstructions. In addition, for Experiment 1, the peak concentration in the reconstruction occurs at a point between the two point sampling locations closest to the SF_6 source. Perhaps if we had used a higher point sampling density near the source, we would have measured a peak point concentration closer to that predicted by the reconstruction. Given the excellent reconstruction of the peak in the kriged data (comparing Figs 3e and b), we do not believe that any feature of the SBFM significantly distorted the experimental peak.

Figure 4f shows a markedly flatter profile with a lower peak concentration than Figs 4b and e. Again, we believe this reflects the nature of the measured ray-integral data rather than an artifact of the SBFM routine. When comparing the ray integrals generated from the kriged synthetic data to those measured by the FTIR, we find that in the synthetic ray-integral data the differences among rays are far more pronounced than in the experimental data: the highest synthetic ray-integral concentrations are up to 30% higher than the measured ones and the lowest synthetic concentrations are lower than the lowest measured ones. This observation may be partially due to vertical concentration gradients. The cross-section of the FTIR beam was approximately 20 cm, and so it sampled 20 cm in the vertical as well as the horizontal direction. The point concentration measurements were sampled 5 cm above the lower edge of the FTIR beam. Thus, most of the volume scanned by the FTIR was slightly above the plane of the point samplers. Since the SF_6 source was beneath the plane of measurement and the air exhaust was on the ceiling, we expect concentration gradients in the horizontal plane to decrease with height. This would cause the FTIR to sample a more even concentration distribution than the point samplers.

Temporal variability during Experiment 2 probably also contributed to the discrepancy between Fig. 4f and Fig. 4a/b. The four highest concentrations shown in Fig. 4a were derived from the time-resolved B&K data. The SF6 concentration at these points exhibited significant variability during the experiments: at two of these points, concentrations steadily decreased from $\sim$ 14 ppm to $\sim$ 0.5 ppm; at a third point, the concentration steadily increased from $\sim$ 1 ppm to $\sim$ 12 ppm; and at the fourth point, the concentration fluctuated between 4 and 10 ppm. We believe that these changes occurred owing to varying temperatures inside and outside the building, and the resulting changes in the mixed convection flow in the enclosure. Thus, the highest point concentrations plotted in Figs 4a and b do not represent truly time-stationary concentrations, and therefore perfect agreement between Fig. 4a/b and Fig. 4f should not be expected.

The σ_{ray} values for Figs 3f and 4f are 1.3 and four times as high, respectively, as the same values for Figs 3e and 4e. There are probably several reasons for this (these reasons also apply to the ART reconstructions). The most obvious is that the former reconstructions are based on experimental data, which have errors associated with them, while the latter are based on internally consistent synthetic data. The error in the experimental data comes partly from the FTIR and partly from the sequential sampling and time-averaging process associated with repeated, non-instantaneous scans. One would prefer to reconstruct concentrations based on snapshots of rays throughout the plane (i.e. all ray-integral concentrations would be sampled simultaneously). Such data would not suffer from any temporal fluctuations in the concentration distribution which lead to inconsistencies in the ray-integral data set. Unfortunately, given technological constraints, this objective is prohibitively costly at present.

Other factors contributing to the inferior σ_{ray} values for the reconstructions of the experimental ray-inte-

gral data are the artificial discretization of the reconstruction plane into pixels in ART and the assumption of constant concentration along the ray width in SBFM. These factors create a discrepancy between the measured path integration, which is truly continuous, and the computed integration. For regions of low spatial variability in concentration, this discrepancy is likely to be of minor consequence; however, it can have a marked effect in areas of steep concentration gradients. We have observed this effect with SBFM. At the outset of this work, SBFM was applied to the experimental data while still employing a pixel discretization. By switching to line integration, agreement between measured and reconstructed ray integrals was greatly improved for rays passing through areas of steep concentration gradients (e.g. close to the dominant peak in Fig. 3b).

Choosing the proper number of Gaussians

What is the appropriate number of Gaussian basis functions for modeling the data in a CT reconstruction? Clearly, one Gaussian would be the minimum necessary to represent anything but a perfectly mixed condition. At the other extreme, using as many Gaussians as pixels would constitute a functional equivalent of an ART reconstruction. Between these limits, it seems that the number of Gaussian basis functions should be guided by the principle of Occam's razor: parameters introduced to the model must not be increased beyond necessity. Evaluating the χ_ν^2 of a reconstruction helps to determine when this point has been reached, because χ_ν^2 is the ratio of the estimated variance of the fit to the variance of the parent distribution, normalized by the number of degrees of freedom n. If the fitting function is a good approximation to the parent distribution and the uncertainties in the measured parameters are correctly quantified, then χ_ν^2 should be approximately unity.

Figure 5 shows how χ_ν^2 and σ_{ray} change with the number of fitting parameters for the two experiments. The first point in each graph represents the case in which the concentration distribution is modeled as a constant. Thereafter, there are six fitting parameters for every Gaussian added to the model. As expected, σ_{ray} decreases each time the number of fitting parameters is increased, but the decrease diminishes for each successive case. χ_ν^2 mostly decreases with σ_{ray}, but may increase slightly if σ_{ray} decreases only very slightly (Fig. 5b).

For Experiment 1, the value of χ_ν^2 decreases below 2, to 1.2, when using three Gaussians, indicating that three Gaussians should be an appropriate number. It is useful also to look at the behavior of σ_{ray}, in which all rays carry equal weight. Between a constant concentration and one Gaussian, σ_{ray} hardly decreases: the single Gaussian produced by the minimization is broad and flat, thus behaving almost like a constant concentration. Most ray integrals are hardly affected by this relatively minor change in the concentration

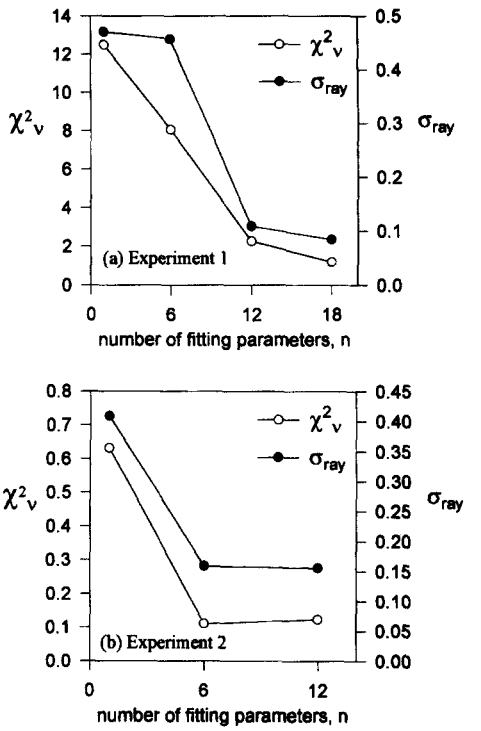

Fig. 5. Effect of number of fitting parameters on quality of fit for SBFM reconstruction based on measured ray-integral data: (a) Experiment 1; (b) Experiment 2. See equations (3) and (4) for parameter definition.

profile. With two Gaussians, a dominant peak appears. This allows for substantial improvement in the fit for rays passing through the peak, which is reflected in a substantial decrease in σ_{ray}. Adding a third Gaussian allows the secondary concentration peak to be refined, leading to a further significant (20%) decrease in σ_{ray}. An additional Gaussian would not be expected to lead to a further significant decrease, as both σ_{ray} and χ_ν^2 are leveling out at this point.

Since the χ_ν^2 value for the SBFM reconstruction of Experiment 1 is larger than unity, the overall difference between the ray integrals of the fit and the measured ray integrals is not within the estimated measurement error. Two likely explanations exist: (1) the true concentration profile is not well represented by bivariate Gaussians and a different functional form would be more appropriate; or (2) the true uncertainty in the measured ray integrals is larger than that estimated by using σ_r^2. The latter factor is likely to be a significant contributor. Drift of the SF_6 plume during the experiment or a change in the air flow pattern within the room due to temperature changes, and thus a drift in the ray-integral concentration of some of the rays, would lead to an underestimation of σ_r^2 and thereby a nonideal χ_ν^2 value. In fact, there is evidence of such drift in the time-resolved ray-integral data: sets of several adjacent rays manifest similar patterns of change from one sweep of the FTIR to the next. Hence, path-integral concentration values for each

ray are not necessarily Gaussian distributed, thereby violating an underlying assumption of chi-squared theory. Nevertheless, the chi-squared approach provides a useful framework for interpreting reconstruction results.

For Experiment 2, both χ_ν^2 and σ_{ray} change very slightly in going from one to two Gaussians, so it seems that a single Gaussian is adequate. The values of χ_ν^2 are all about one tenth of what we would expect (the best value is around 0.11 instead of unity). Most of the ray variances were rather large in this experiment. Whereas the mean relative standard deviation of the time-averaged ray-integral concentrations was 15% in Experiment 1, it was 79% in Experiment 2. The quality of the fit, as quantified by σ_{ray} for the final result, is also worse (0.084 for Experiment 1, 0.16 for Experiment 2). The large relative standard deviations in Experiment 2 indicate strong temporal fluctuations in concentration during that experiment, causing the mean path-integral concentration for each ray to be less accurate than in Experiment 1. With the ray variances in Experiment 2 being disproportionately larger, however, the χ_ν^2 value was correspondingly smaller for this experiment. As discussed in the previous paragraph, the χ_ν^2 value alone is therefore not enough to judge the goodness of fit. It must be analyzed in conjunction with the trend in σ_{ray} and detailed characteristics of the input data.

CONCLUSIONS

Our investigation into measuring airborne tracer gas concentrations using optical remote sensing and computed tomography has led to the following conclusions regarding reconstruction methods:

(1) Iterative algebraic reconstruction techniques (ART), when used with a ray geometry designed for rapid scanning, result in concentration profiles characterized by unrealistically high spatial variability. Deceptively low values of σ_{ray} make it clear that applying ART to the reconstruction of gas concentrations in air using such ray geometries could lead to unrecognized, significant errors in estimates of the gas concentration profiles.

(2) The new approach we have developed, of modeling the concentration profile with smooth basis functions (SBFM), produces reconstructions that are much more consistent than ART when compared with measured point-sample concentration data. Bivariate Gaussians constitute a satisfactory basis for representing the concentration profiles generated in our ventilated chamber experiments.

(3) The SBFM approach obviates the need for a trade-off between spatial resolution and the number of free parameters in the reconstruction, a critical weakness of ART, where the number of pixels determines both the level of spatial resolution and the number of free parameters.

(4) The minimization of σ_{ray}, when used in conjunction with χ_ν^2 values, provides a basis for determining the proper number of parameters necessary to represent the data in a manner consistent with the amount of information provided by the rays.

The SBFM routine using analytical integration and executed on a 486 PC (33 SX) requires about 2 h for a comprehensive search (annealing) for the global minimum of the goodness-of-fit parameter σ_{ray} from a random first guess using two Gaussians. The computing time increases with the number of Gaussians. If one can make a good first guess at the approximate location and height of peaks, the annealing need not be quite as comprehensive and the computing time can be significantly shortened. A first guess might be generated from the trends visible in an ART reconstruction, which can be completed in a matter of minutes. Another option would be to augment the SBFM program with a routine that searches for areas in the reconstruction plane through which many rays with high integral concentration pass, and placing the maxima of the Gaussians there for the first guess. The minimization might also be substantially accelerated by improving the program, including changing to a different programming language, such as C. It is possible that with a more powerful computer, near real-time monitoring, one goal for CT-FTIR development, could be achieved with SBFM.

Rapid progress has been made recently both in remote sensing technologies and in understanding the characteristics of CT reconstruction algorithms for air quality applications. A significant effort remains to demonstrate in a broad range of field tests the circumstances under which coupling CT with remote sensing can yield accurate measurement results. Particular attention is warranted to the problem of reconstructing time-varying concentration profiles. Also, the design of practical monitoring systems for emissions characterization in outdoor air represents a worthwhile, yet unrealized objective. Although still in the early stages of development, the use of computed tomography in conjunction with remote sensing holds substantial ultimate promise for fast and cost-effective gas concentration distribution measurements over a range of spatial scales, with potential application to many indoor and outdoor air quality problems.

Acknowledgements—We would like to thank Marc Fischer for many fruitful discussions and helpful suggestions in developing the SBFM method. The experiments were performed in collaboration with Doo Yong Park from the University of Michigan, with resource support from Michael Yost (University of Washington), Steven Levine (University of Michigan), and Robert Spear (University of California). We also gratefully acknowledge the support of MDA Scientific (now Environmental Technology Group (ETG)), which provided the FTIR and the retroreflectors. This work was supported by NIOSH grant 1R01-OH 02666, by National Science Foundation grant BCS-9057298 and by the Assistant Secretary for Conservation and Renewable

Energy, Office of Building Technologies, Building Systems Division of the U.S. Department of Energy under Contract no. DE-AC03-76SF00098.

REFERENCES

Bevington P. R. (1969) *Data Reduction and Error Analysis for the Physical Sciences*, Chapter 10. McGraw-Hill, New York.

Brooks R. A. and DiChiro G. (1975) Theory of image reconstruction in computed tomography. *Radiology* **117**, 561–572.

Drescher A. C., Park D. Y., Yost M. G., Gadgil A. J., Levine S. P. and Nazaroff W. W. (1995) Stationary and time-dependent indoor tracer-gas concentration profiles measured by OP-FTIR remote sensing and SBFM computed tomography. *Atmospheric Environment*, submitted.

Faris G. W. and Byer R. L. (1987) Beam deflection optical tomography of a flame. *Optics Lett.* **12**, 155–157.

Fischer M. L., Drescher A. C., Gadgil A. J. and Yost M. G. (1994) A system for rapid detection and mapping of gas plumes on 100 m scales: examination of some technical and economic issues. In *Proc. International Symp. on Optical Remote Sensing for Environmental and Process Monitoring*. 7–10 November, 1994, McLean, VA, Air and Waste Management Assoc.

Jones T. A., Hamilton D. E. and Johnson C. R. (1986) *Contouring Geologic Surfaces with the Computer*. Van Nostrand Reinhold, New York.

Press W. H., Teukolsky S. A., Vetterling W. T. and Flannery B. P. (1992) *Numerical Recipes in C*, 2nd Ed. Cambridge University Press, Cambridge, MA.

Todd L. A. (1992a) Optical remote sensing/computed tomography systems for workplace exposure assessments. In *Proc. Optical Remote Sensing Applications to Environmental and Industrial Safety Problems*, pp. 356–360. Air and Waste Manage. Assoc. SP81.

Todd L. A. (1992b) Optical remote sensing/computed tomography beam geometries for monitoring workplace gases and vapors. In *Proc. Optical Remote Sensing Applications to Environmental and Industrial Safety Problems*, pp. 390–393. Air and Waste Manage. Assoc. SP81.

Todd L. and Leith D. (1990) Remote sensing and computed tomography in industrial hygiene. *Am. Ind. Hyg. Ass. J.* **51**, 224–233.

Todd L. and Ramachandran N. (1994) Evaluation of algorithms for tomographic reconstruction of chemical concentrations in indoor air. *Am. Ind. Hyg. Ass. J.* **55**, 403–417.

Wolfe D. C. (1980) On the application of optical computed tomography to remote air pollution measurements. Ph.D. dissertation, Stanford University, Palo Alto, CA.

Yost M. G., Gadgil A. J., Drescher A. C., Zhou Y., Simonds M. A., Levine S. P., Nazaroff W. W. and Saisan P. A. (1994) Imaging indoor tracer-gas concentrations with computed tomography: experimental results with a remote sensing FTIR system. *Am. Ind. Hyg. Assoc. J.* **55**, 395–402.